

(10) **Patent No.:** **US 9,262,343 B2**  
(45) **Date of Patent:** **\*Feb. 16, 2016**

USPC ..... 711/118, 125, E12.008, 163  
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,586,297	A	12/1996	Bryg et al.	
6,349,361	B1	2/2002	Altman et al.	
7,552,290	B2 *	6/2009	Alsup .....	711/152
7,966,459	B2	6/2011	Nussbaum et al.	
8,291,379	B2	10/2012	Krauss et al.	
8,327,188	B2	12/2012	Karlsson et al.	
8,356,166	B2 *	1/2013	Adl-Tabatabai et al. ....	712/229
8,688,920	B2 *	4/2014	Heller, Jr. ....	711/150
2004/0044850	A1	3/2004	George et al.	
2010/0023707	A1	1/2010	Hohmuth et al.	
2010/0205408	A1 *	8/2010	Chung .....	G06F 9/466 712/216

2004/0044850	A1	3/2004	George et al.	
2010/0023707	A1	1/2010	Hohmuth et al.	
2010/0205408	A1*	8/2010	Chung .....	G06F 9/466 712/216

2011/0125973	A1	5/2011	Lev et al.
2011/0208921	A1	8/2011	Pohlack et al.
2011/0246724	A1	10/2011	Marathe et al.
2011/0296114	A1	12/2011	Farrell et al.
2012/0005461	A1	1/2012	Moir et al.

(Continued)

## OTHER PUBLICATIONS

Pending U.S. Appl. No. 14/225,482 , “Transactional Processing Based Upon Run-Time Storage Values ”, filed Mar. 26, 2014.

(Continued)

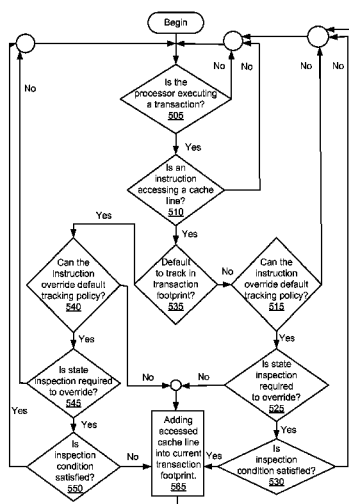
*Primary Examiner* — Pierre-Michel Bataille

(74) *Attorney, Agent, or Firm* — Teddi E. Maranzano;  
William Kinnaman

(57) **ABSTRACT**

A memory access operand of an instruction that accesses memory may be treated as a transaction atomic access. The processor may execute one or more processor state setting instructions, causing state information to be set in the processor. Upon executing a transaction policy override instruction, the default conflict detection policy is overridden for one or more subsequent memory accessing instructions.

**20 Claims, 6 Drawing Sheets**



(56)

**References Cited**

## U.S. PATENT DOCUMENTS

2012/0110317 A1\* 5/2012 Scheer ..... G06F 21/10  
713/150  
2012/0117334 A1 5/2012 Sheaffer et al.  
2012/0278120 A1\* 11/2012 Insko ..... G06F 9/4881  
705/7.21  
2014/0379996 A1\* 12/2014 Rajwar et al. .... 711/145

## OTHER PUBLICATIONS

Proceedings 2012 IEEE/ACM 45th International Symposium on  
Microarchitecture MICRO-45, presented Dec. 1-5, 2012, "Transac-  
tional Memory Architecture and Implementation for IBM System z",  
pp. 25-36.

IBM, "Principles of Operation", Tenth Edition (Sep. 2012), SA22-  
7832-09.

"Intel® Architecture Instruction Set Extensions Programming Ref-  
erence", 319433-012A, Feb. 2012.

Austen McDonald, "Architectures for Transactional Memory", A  
Dissertation Submitted to the Department of Computer Science and  
the Committee on Graduate Studies of Stanford University in Partial  
Fulfillment of the Requirements for the Degree of Doctor of Philoso-  
phy, Jun. 2009, pp. 1-145.

P. Mak et al., IBM J. Res. & Dev. vol. 53, No. 1, Paper 2, 2009, "IBM  
System z10 Processor Cache Subsystem Microarchitecture", pp. 2:1-  
2:12.

Mentis et al.; "Model Checking and Code Generation for Transac-  
tional Processing Software", Concurrency and Comp. Pra. & Exp.  
Journal, vol. 24, Issue 7, May 2012, pp. 1-17.

Chung et al.; "Analysis on Semantic Transactional Memory Footprint  
for Hardware Transactional Memory", IISWC IEEE Inter. Conf. on,  
Dec. 2-4, 2010, pp. 1-11.

\* cited by examiner




FIG. 1

FIG. 2

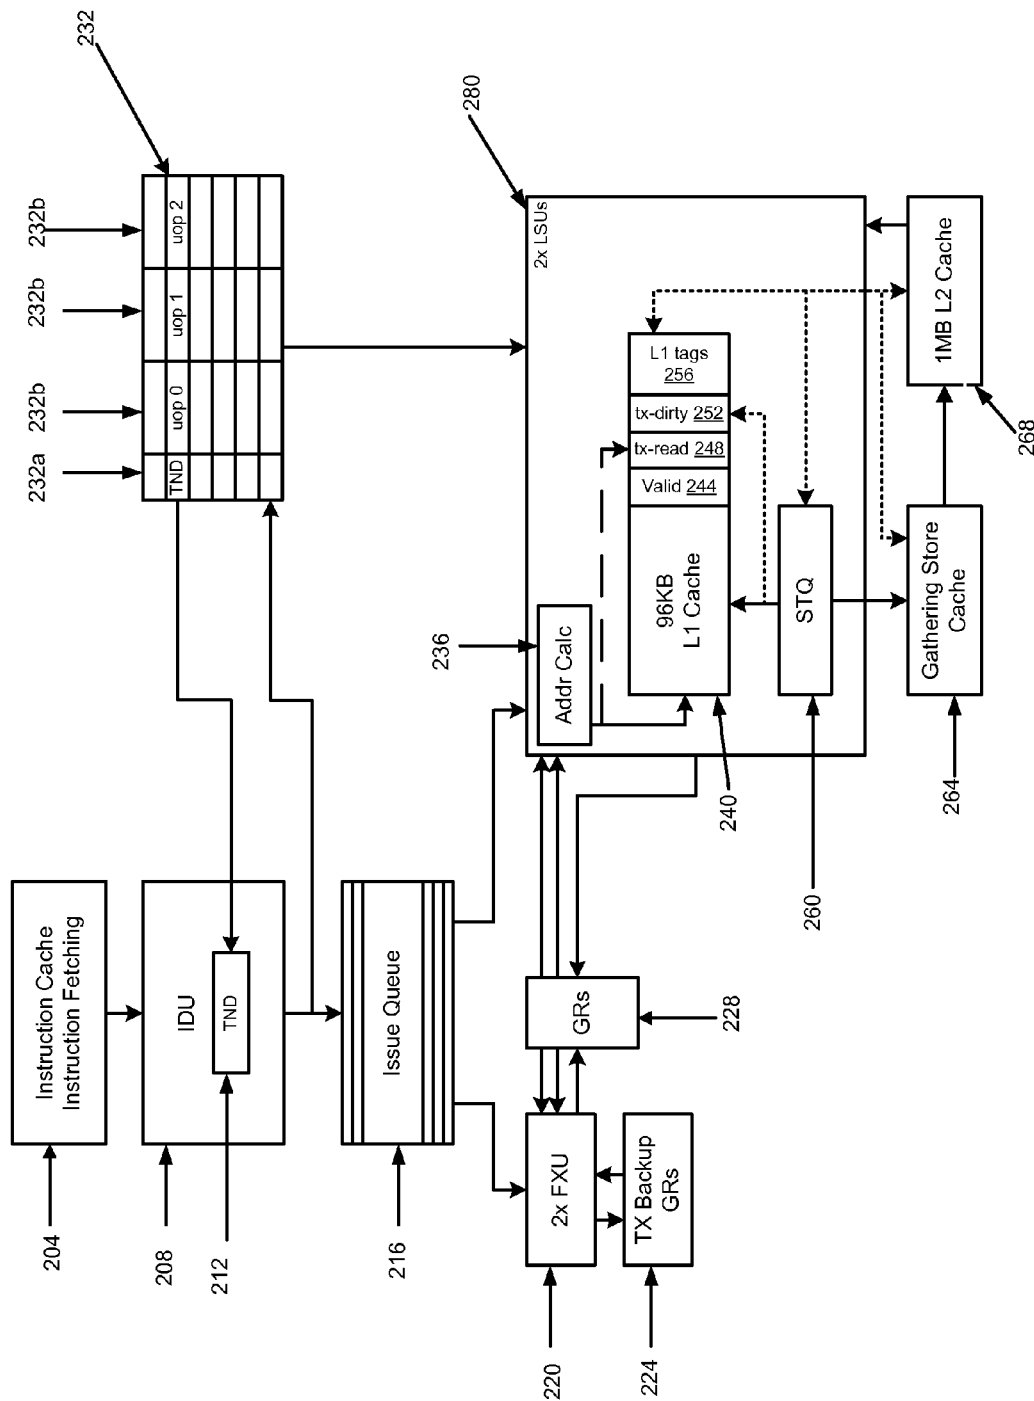


FIG. 3

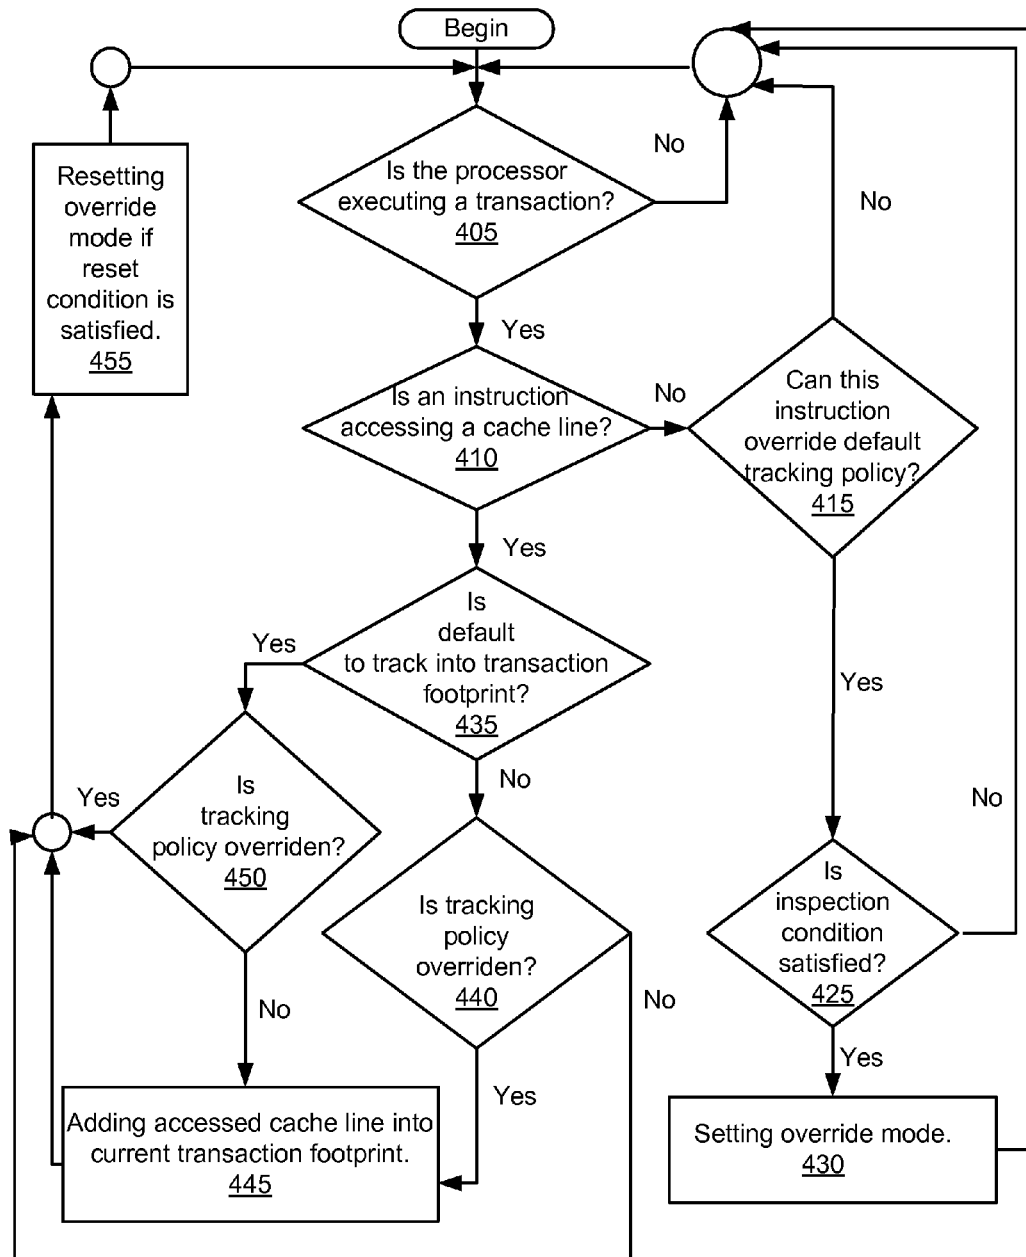


FIG. 4

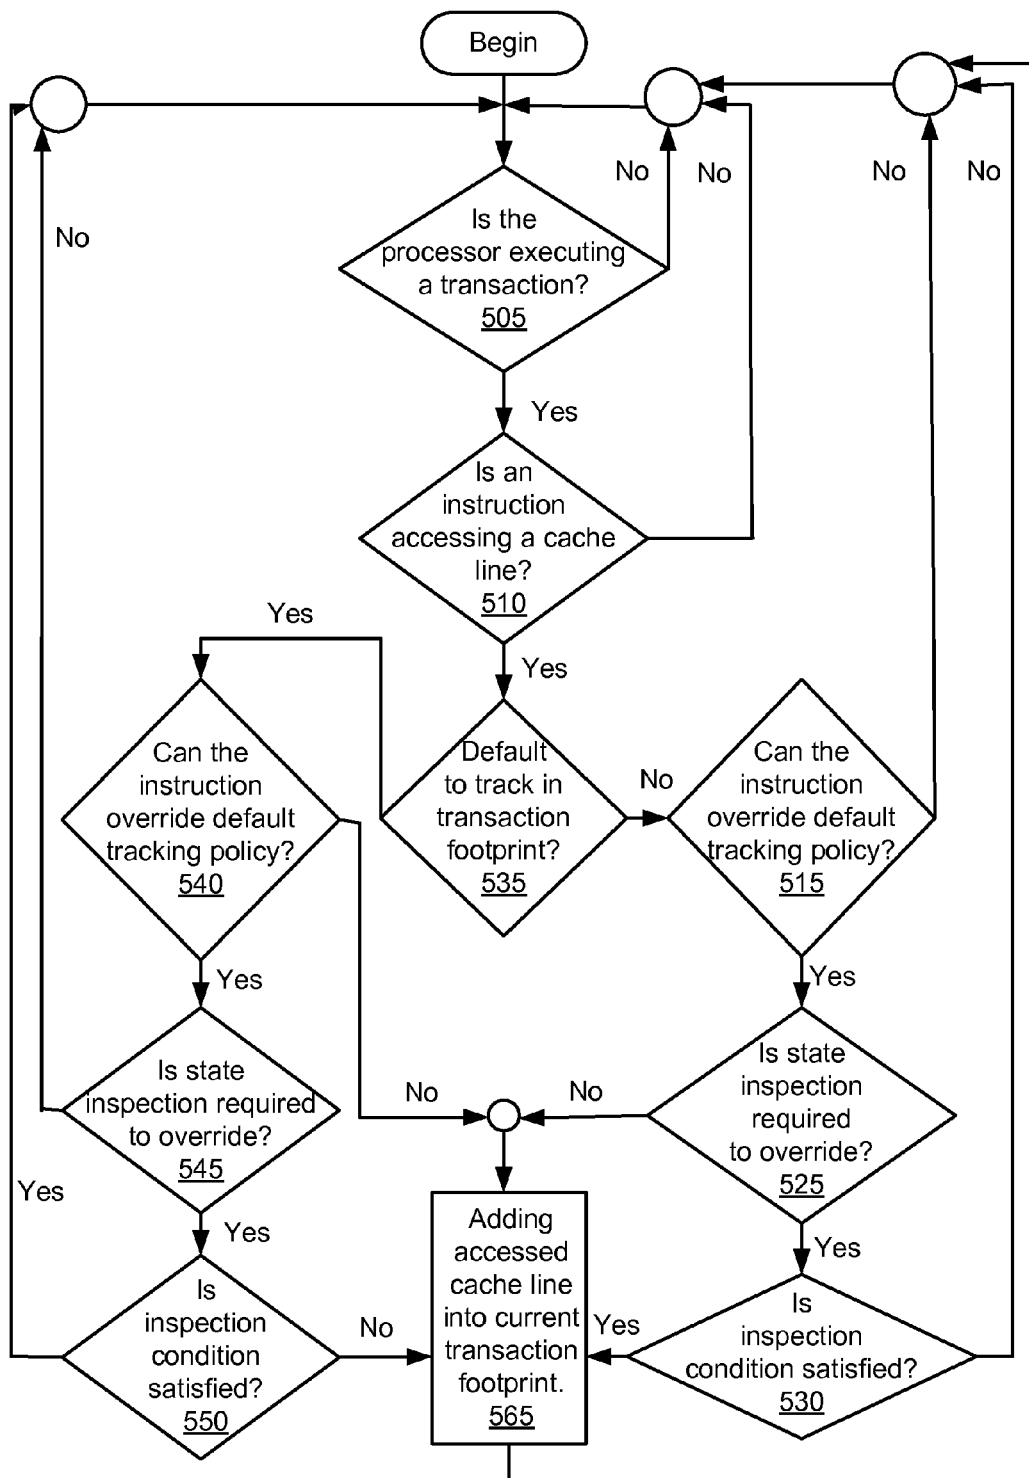


FIG. 5

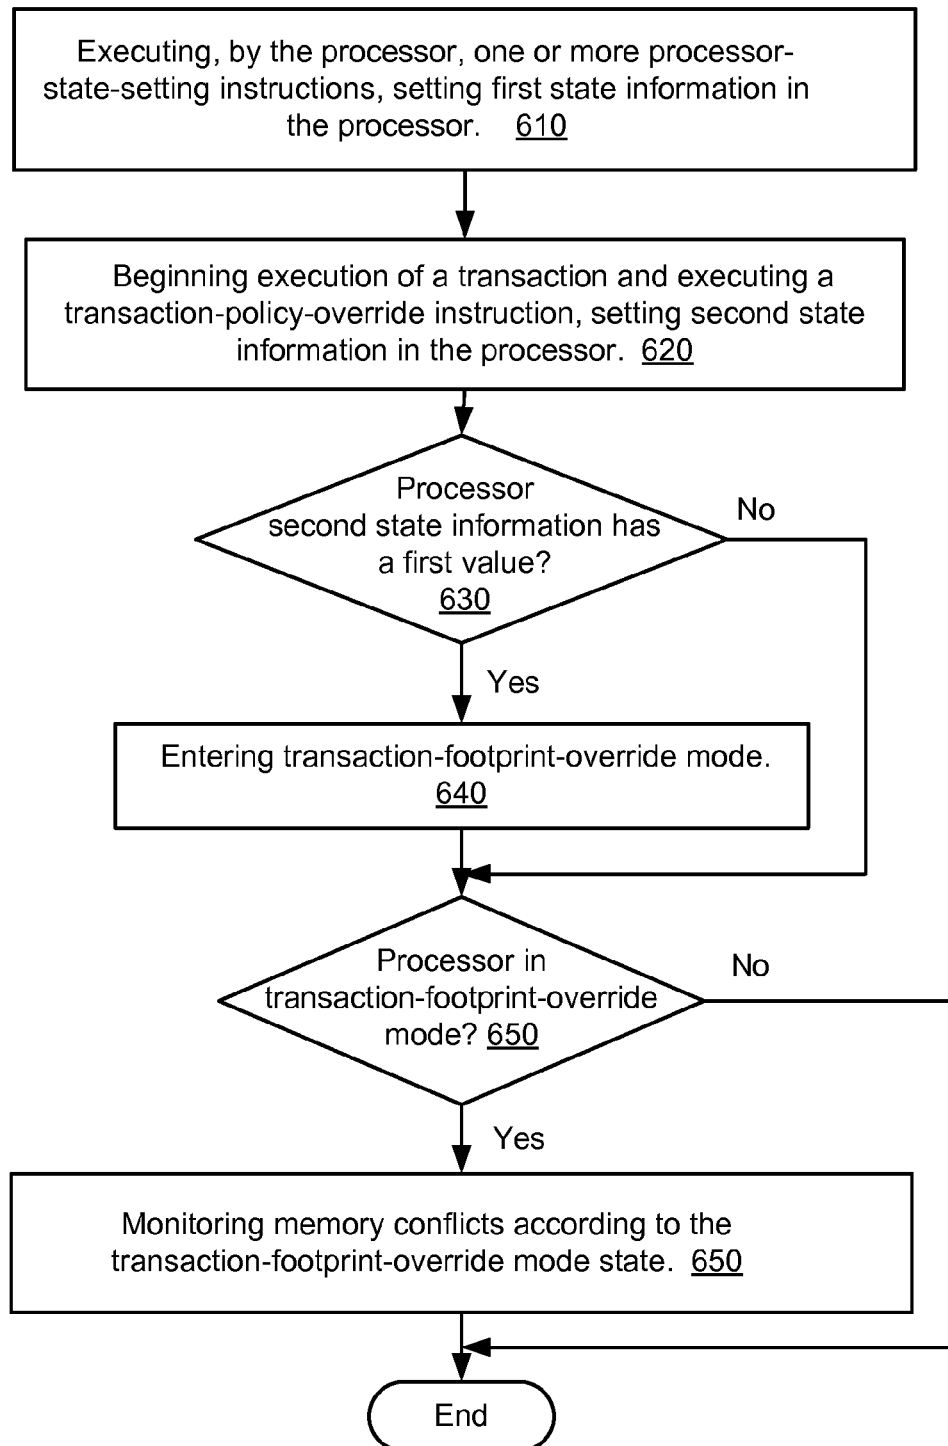


FIG. 6




FIG. 7



1

## TRANSACTIONAL PROCESSING BASED UPON RUN-TIME CONDITIONS

### FIELD

This disclosure relates generally to transactional memory systems and more specifically to a method, computer program and computer system for transactional processing.

### BACKGROUND

The number of central processing unit (CPU) cores on a chip and the number of CPU cores connected to a shared memory continues to grow significantly to support growing workload capacity demand. The increasing number of CPUs cooperating to process the same workloads puts a significant burden on software scalability; for example, shared queues or data-structures protected by traditional semaphores become hot spots and lead to sub-linear n-way scaling curves. Traditionally this has been countered by implementing finer-grained locking in software, and with lower latency/higher bandwidth interconnects in hardware. Implementing fine-grained locking to improve software scalability can be very complicated and error-prone, and at today's CPU frequencies, the latencies of hardware interconnects are limited by the physical dimension of the chips and systems, and by the speed of light.

Implementations of hardware Transactional Memory (HTM, or in this discussion, simply TM) have been introduced, wherein a group of instructions—called a transaction—operate in an atomic manner on a data structure in memory, as viewed by other central processing units (CPUs) and the I/O subsystem (atomic operation is also known as “block concurrent” or “serialized” in other literature). The transaction executes optimistically without obtaining a lock, but may need to abort and retry the transaction execution if an operation, of the executing transaction, on a memory location conflicts with another operation on the same memory location. Previously, software transactional memory implementations have been proposed to support software Transactional Memory (TM). However, hardware TM can provide improved performance aspects and ease of use over software TM.

U.S. Patent Application Publication No 2004/0044850 titled “Method and apparatus for the synchronization of distributed caches” filed 2002 Aug. 28 and incorporated by reference herein teaches a method and apparatus for the synchronization of distributed caches. More particularly, the present embodiment to cache memory systems and more particularly to a hierarchical caching protocol suitable for use with distributed caches, including use within a caching input/output (I/O) hub.

U.S. Pat. No. 5,586,297 titled “Partial cache line write transactions in a computing system with a write back cache” filed 1994 Mar. 24 and incorporated by reference herein teaches a computing system is presented which includes a memory, an input/output adapter and a processor. The processor includes a write back cache in which dirty data may be stored. When performing a coherent write from the input/output adapter to the memory, a block of data is written from the input/output adapter to a memory location within the memory. The block of data contains less data than a full cache line in the write back cache. The write back cache is searched to determine whether the write back cache contains data for the memory location. When the search determines that the

2

write back cache contains data for the memory location a full cache line which contains the data for the memory location is purged.

### BRIEF SUMMARY

Dynamically adjusting a policy of whether a memory operation and access of a memory accessing instruction is to be treated as a transaction atomic access in a transactional execution environment is provided. An exemplary embodiment of the disclosure may include executing, by the processor, one or more processor-state-setting instructions, the executing causing processor first state information to be set in the processor, the processor first state information defined by an instruction set architecture (ISA) of the processor; and beginning execution of a transaction, the execution comprising: executing, by the processor, a transaction-policy-override instructions, the executing causing processor second state information to be set in the processor, the processor second state information defined by the ISA of the processor; based on the processor second state information having a first value, entering a transaction-footprint-override mode; and based on the processor being in transaction-footprint-override mode, any one of ignoring memory conflicts that would otherwise be monitored and monitoring memory conflicts that would otherwise be ignored.

### BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWINGS

Features and advantages of the present disclosure will become apparent from the following detailed description of illustrative embodiments thereof, which is to be read in connection with the accompanying drawings. The various features of the drawings are not to scale as the illustrations are for clarity in facilitating one skilled in the art in understanding the disclosure in conjunction with the detailed description. In the drawings:

FIGS. 1 and 2 depict an example multicore Transactional Memory environment, in accordance with embodiments of the present disclosure;

FIG. 3 depicts example components of an example CPU, in accordance with embodiments of the present disclosure;

FIG. 4 depicts a flow diagram of a method for processing one or more instructions in a transactions based upon run-time conditions, according to an exemplary embodiment;

FIG. 5 depicts a flow diagram of a method for processing an instruction in a transactions based upon run-time conditions, according to an exemplary embodiment;

FIG. 6 depicts an exemplary flow of transactional processing; and

FIG. 7 is a schematic block diagram of hardware and software of the computer environment according to at least one exemplary embodiment of the method of FIGS. 4-5.

### DETAILED DESCRIPTION

Historically, a computer system or processor had only a single processor (aka processing unit or central processing unit). The processor included an instruction processing unit (IPU), a branch unit, a memory control unit and the like. Such processors were capable of executing a single thread of a program at a time. Operating systems were developed that could time-share a processor by dispatching a program to be executed on the processor for a period of time, and then dispatching another program to be executed on the processor for another period of time. As technology evolved, memory

subsystem caches were often added to the processor as well as complex dynamic address translation including translation lookaside buffers (TLBs). The IPU itself was often referred to as a processor. As technology continued to evolve, an entire processor could be packaged as a single semiconductor chip or die, such a processor was referred to as a microprocessor. Then processors were developed that incorporated multiple IPU's, such processors were often referred to as multi-processors. Each such processor of a multi-processor computer system (processor) may include individual or shared caches, memory interfaces, system bus, address translation mechanism and the like. Virtual machine and instruction set architecture (ISA) emulators added a layer of software to a processor, that provided the virtual machine with multiple "virtual processors" (aka processors) by time-slice usage of a single IPU in a single hardware processor. As technology further evolved, multi-threaded processors were developed, enabling a single hardware processor having a single multi-thread IPU to provide a capability of simultaneously executing threads of different programs, thus each thread of a multi-threaded processor appeared to the operating system as a processor. As technology further evolved, it was possible to put multiple processors (each having an IPU) on a single semiconductor chip or die. These processors were referred to processor cores or just cores. Thus the terms such as processor, central processing unit, processing unit, microprocessor, core, processor core, processor thread, and thread, for example, are often used interchangeably. Aspects of embodiments herein may be practiced by any or all processors including those shown supra, without departing from the teachings herein. Wherein the term "thread" or "processor thread" is used herein, it is expected that particular advantage of the embodiment may be had in a processor thread implementation.

#### Transaction Execution in Intel® Based Embodiments

In "Intel® Architecture Instruction Set Extensions Programming Reference" 319433-012A, February 2012, incorporated herein by reference in its entirety, Chapter 8 teaches, in part, that multithreaded applications may take advantage of increasing numbers of CPU cores to achieve higher performance. However, the writing of multi-threaded applications requires programmers to understand and take into account data sharing among the multiple threads. Access to shared data typically requires synchronization mechanisms. These synchronization mechanisms are used to ensure that multiple threads update shared data by serializing operations that are applied to the shared data, often through the use of a critical section that is protected by a lock. Since serialization limits concurrency, programmers try to limit the overhead due to synchronization.

Intel® Transactional Synchronization Extensions (Intel® TSX) allow a processor to dynamically determine whether threads need to be serialized through lock-protected critical sections, and to perform that serialization only when required. This allows the processor to expose and exploit concurrency that is hidden in an application because of dynamically unnecessary synchronization.

With Intel® TSX, programmer-specified code regions (also referred to as "transactional regions" or just "transactions") are executed transactionally. If the transactional execution completes successfully, then all memory operations performed within the transactional region will appear to have occurred instantaneously when viewed from other processors. A processor makes the memory operations of the executed transaction, performed within the transactional region, visible to other processors only when a successful

commit occurs, i.e., when the transaction successfully completes execution. This process is often referred to as an atomic commit.

Intel® TSX provides two software interfaces to specify regions of code for transactional execution. Hardware Lock Elision (HLE) is a legacy compatible instruction set extension (comprising the XACQUIRE and XRELEASE prefixes) to specify transactional regions. Restricted Transactional Memory (RTM) is a new instruction set interface (comprising the XBEGIN, XEND, and XABORT instructions) for programmers to define transactional regions in a more flexible manner than that possible with HLE. HLE is for programmers who prefer the backward compatibility of the conventional mutual exclusion programming model and would like to run HLE-enabled software on legacy hardware but would also like to take advantage of the new lock elision capabilities on hardware with HLE support. RTM is for programmers who prefer a flexible interface to the transactional execution hardware. In addition, Intel® TSX also provides an XTEST instruction. This instruction allows software to query whether the logical processor is transactionally executing in a transactional region identified by either HLE or RTM.

Since a successful transactional execution ensures an atomic commit, the processor executes the code region optimistically without explicit synchronization. If synchronization was unnecessary for that specific execution, execution can commit without any cross-thread serialization. If the processor cannot commit atomically, then the optimistic execution fails. When this happens, the processor will roll back the execution, a process referred to as a transactional abort. On a transactional abort, the processor will discard all updates performed in the memory region used by the transaction, restore architectural state to appear as if the optimistic execution never occurred, and resume execution non-transactionally.

A processor can perform a transactional abort for numerous reasons. A primary reason to abort a transaction is due to conflicting memory accesses between the transactionally executing logical processor and another logical processor. Such conflicting memory accesses may prevent a successful transactional execution. Memory addresses read from within a transactional region constitute the read-set of the transactional region and addresses written to within the transactional region constitute the write-set of the transactional region. Intel® TSX maintains the read- and write-sets at the granularity of a cache line. A conflicting memory access occurs if another logical processor either reads a location that is part of the transactional region's write-set or writes a location that is a part of either the read- or write-set of the transactional region. A conflicting access typically means that serialization is required for this code region. Since Intel® TSX detects data conflicts at the granularity of a cache line, unrelated data locations placed in the same cache line will be detected as conflicts that result in transactional aborts. Transactional aborts may also occur due to limited transactional resources. For example, the amount of data accessed in the region may exceed an implementation-specific capacity. Additionally, some instructions and system events may cause transactional aborts. Frequent transactional aborts result in wasted cycles and increased inefficiency.

#### Hardware Lock Elision

Hardware Lock Elision (HLE) provides a legacy compatible instruction set interface for programmers to use transactional execution. HLE provides two new instruction prefix hints: XACQUIRE and XRELEASE.

With HLE, a programmer adds the XACQUIRE prefix to the front of the instruction that is used to acquire the lock that

is protecting the critical section. The processor treats the prefix as a hint to elide the write associated with the lock acquire operation. Even though the lock acquire has an associated write operation to the lock, the processor does not add the address of the lock to the transactional region's write-set nor does it issue any write requests to the lock. Instead, the address of the lock is added to the read-set. The logical processor enters transactional execution. If the lock was available before the XACQUIRE prefixed instruction, then all other processors will continue to see the lock as available afterwards. Since the transactionally executing logical processor neither added the address of the lock to its write-set nor performed externally visible write operations to the lock, other logical processors can read the lock without causing a data conflict. This allows other logical processors to also enter and concurrently execute the critical section protected by the lock. The processor automatically detects any data conflicts that occur during the transactional execution and will perform a transactional abort if necessary.

Even though the eliding processor did not perform any external write operations to the lock, the hardware ensures program order of operations on the lock. If the eliding processor itself reads the value of the lock in the critical section, it will appear as if the processor had acquired the lock, i.e. the read will return the non-elided value. This behavior allows an HLE execution to be functionally equivalent to an execution without the HLE prefixes.

An XRELEASE prefix can be added in front of an instruction that is used to release the lock protecting a critical section. Releasing the lock involves a write to the lock. If the instruction is to restore the value of the lock to the value the lock had prior to the XACQUIRE prefixed lock acquire operation on the same lock, then the processor elides the external write request associated with the release of the lock and does not add the address of the lock to the write-set. The processor then attempts to commit the transactional execution.

With HLE, if multiple threads execute critical sections protected by the same lock but they do not perform any conflicting operations on each other's data, then the threads can execute concurrently and without serialization. Even though the software uses lock acquisition operations on a common lock, the hardware recognizes this, elides the lock, and executes the critical sections on the two threads without requiring any communication through the lock—if such communication was dynamically unnecessary.

If the processor is unable to execute the region transactionally, then the processor will execute the region non-transactionally and without elision. HLE enabled software has the same forward progress guarantees as the underlying non-HLE lock-based execution. For successful HLE execution, the lock and the critical section code must follow certain guidelines. These guidelines only affect performance; and failure to follow these guidelines will not result in a functional failure. Hardware without HLE support will ignore the XACQUIRE and XRELEASE prefix hints and will not perform any elision since these prefixes correspond to the REPNE/REPE IA-32 prefixes which are ignored on the instructions where XACQUIRE and XRELEASE are valid. Importantly, HLE is compatible with the existing lock-based programming model. Improper use of hints will not cause functional bugs though it may expose latent bugs already in the code.

Restricted Transactional Memory (RTM) provides a flexible software interface for transactional execution. RTM provides three new instructions—XBEGIN, XEND, and XABORT—for programmers to start, commit, and abort a transactional execution.

The programmer uses the XBEGIN instruction to specify the start of a transactional code region and the XEND instruction to specify the end of the transactional code region. If the RTM region could not be successfully executed transactionally, then the XBEGIN instruction takes an operand that provides a relative offset to the fallback instruction address.

A processor may abort RTM transactional execution for many reasons. In many instances, the hardware automatically detects transactional abort conditions and restarts execution from the fallback instruction address with the architectural state corresponding to that present at the start of the XBEGIN instruction and the EAX register updated to describe the abort status.

The XABORT instruction allows programmers to abort the execution of an RTM region explicitly. The XABORT instruction takes an 8-bit immediate argument that is loaded into the EAX register and will thus be available to software following an RTM abort. RTM instructions do not have any data memory location associated with them. While the hardware provides no guarantees as to whether an RTM region will ever successfully commit transactionally, most transactions that follow the recommended guidelines are expected to successfully commit transactionally. However, programmers must always provide an alternative code sequence in the fallback path to guarantee forward progress. This may be as simple as acquiring a lock and executing the specified code region non-transactionally. Further, a transaction that always aborts on a given implementation may complete transactionally on a future implementation. Therefore, programmers must ensure the code paths for the transactional region and the alternative code sequence are functionally tested.

#### Detection of HLE Support

A processor supports HLE execution if CPUID.07H.EBX.HLE [bit 4]=1. However, an application can use the HLE prefixes (XACQUIRE and XRELEASE) without checking whether the processor supports HLE. Processors without HLE support ignore these prefixes and will execute the code without entering transactional execution.

#### Detection of RTM Support

A processor supports RTM execution if CPUID.07H.EBX.RTM [bit 11]=1. An application must check if the processor supports RTM before it uses the RTM instructions (XBEGIN, XEND, XABORT). These instructions will generate a #UD exception when used on a processor that does not support RTM.

#### Detection of XTEST Instruction

A processor supports the XTEST instruction if it supports either HLE or RTM. An application must check either of these feature flags before using the XTEST instruction. This instruction will generate a #UD exception when used on a processor that does not support either HLE or RTM.

#### Querying Transactional Execution Status

The XTEST instruction can be used to determine the transactional status of a transactional region specified by HLE or RTM. Note, while the HLE prefixes are ignored on processors that do not support HLE, the XTEST instruction will generate a #UD exception when used on processors that do not support either HLE or RTM.

#### Requirements for HLE Locks

For HLE execution to successfully commit transactionally, the lock must satisfy certain properties and access to the lock must follow certain guidelines.

An XRELEASE prefixed instruction must restore the value of the elided lock to the value it had before the lock acquisition. This allows hardware to safely elide locks by not adding them to the write-set. The data size and data address of the lock release (XRELEASE prefixed) instruction must match

that of the lock acquire (XACQUIRE prefixed) and the lock must not cross a cache line boundary.

Software should not write to the elided lock inside a transactional HLE region with any instruction other than an XRELEASE prefixed instruction, otherwise such a write may cause a transactional abort. In addition, recursive locks (where a thread acquires the same lock multiple times without first releasing the lock) may also cause a transactional abort. Note that software can observe the result of the elided lock acquire inside the critical section. Such a read operation will return the value of the write to the lock.

The processor automatically detects violations to these guidelines, and safely transitions to a non-transactional execution without elision. Since Intel® TSX detects conflicts at the granularity of a cache line, writes to data collocated on the same cache line as the elided lock may be detected as data conflicts by other logical processors eliding the same lock. Transactional Nesting

Both HLE and RTM support nested transactional regions. However, a transactional abort restores state to the operation that started transactional execution: either the outermost XACQUIRE prefixed HLE eligible instruction or the outermost XBEGIN instruction. The processor treats all nested transactions as one transaction.

#### HLE Nesting and Elision

Programmers can nest HLE regions up to an implementation specific depth of MAX\_HLE\_NEST\_COUNT. Each logical processor tracks the nesting count internally but this count is not available to software. An XACQUIRE prefixed HLE-eligible instruction increments the nesting count, and an XRELEASE prefixed HLE-eligible instruction decrements it. The logical processor enters transactional execution when the nesting count goes from zero to one. The logical processor attempts to commit only when the nesting count becomes zero. A transactional abort may occur if the nesting count exceeds MAX\_HLE\_NEST\_COUNT.

In addition to supporting nested HLE regions, the processor can also elide multiple nested locks. The processor tracks a lock for elision beginning with the XACQUIRE prefixed HLE eligible instruction for that lock and ending with the XRELEASE prefixed HLE eligible instruction for that same lock. The processor can, at any one time, track up to a MAX\_HLE\_ELIDED\_LOCKS number of locks. For example, if the implementation supports a MAX\_HLE\_ELIDED\_LOCKS value of two and if the programmer nests three HLE identified critical sections (by performing XACQUIRE prefixed HLE eligible instructions on three distinct locks without performing an intervening XRELEASE prefixed HLE eligible instruction on any one of the locks), then the first two locks will be elided, but the third won't be elided (but will be added to the transaction's write-set). However, the execution will still continue transactionally. Once an XRELEASE for one of the two elided locks is encountered, a subsequent lock acquired through the XACQUIRE prefixed HLE eligible instruction will be elided.

The processor attempts to commit the HLE execution when all elided XACQUIRE and XRELEASE pairs have been matched, the nesting count goes to zero, and the locks have satisfied requirements. If execution cannot commit atomically, then execution transitions to a non-transactional execution without elision as if the first instruction did not have an XACQUIRE prefix.

#### RTM Nesting

Programmers can nest RTM regions up to an implementation specific MAX\_RTM\_NEST\_COUNT. The logical processor tracks the nesting count internally but this count is not available to software. An XBEGIN instruction increments the

nesting count, and an XEND instruction decrements the nesting count. The logical processor attempts to commit only if the nesting count becomes zero. A transactional abort occurs if the nesting count exceeds MAX\_RTM\_NEST\_COUNT.

#### Nesting HLE and RTM

HLE and RTM provide two alternative software interfaces to a common transactional execution capability. Transactional processing behavior is implementation specific when HLE and RTM are nested together, e.g., HLE is inside RTM or RTM is inside HLE. However, in all cases, the implementation will maintain HLE and RTM semantics. An implementation may choose to ignore HLE hints when used inside RTM regions, and may cause a transactional abort when RTM instructions are used inside HLE regions. In the latter case, the transition from transactional to non-transactional execution occurs seamlessly since the processor will re-execute the HLE region without actually doing elision, and then execute the RTM instructions.

#### Abort Status Definition

RTM uses the EAX register to communicate abort status to software. Following an RTM abort the EAX register has the following definition.

TABLE 1

RTM Abort Status Definition		
EAX Register	Bit Position	Meaning
	0	Set if abort caused by XABORT instruction
	1	If set, the transaction may succeed on retry, this bit is always clear if bit 0 is set
	2	Set if another logical processor conflicted with a memory address that was part of the transaction that aborted
	3	Set if an internal buffer overflowed
	4	Set if a debug breakpoint was hit
	5	Set if an abort occurred during execution of a nested transaction
	23:6	Reserved
	31-24	XABORT argument (only valid if bit 0 set, otherwise reserved)

The EAX abort status for RTM only provides causes for aborts. It does not by itself encode whether an abort or commit occurred for the RTM region. The value of EAX can be 0 following an RTM abort. For example, a CPUID instruction when used inside an RTM region causes a transactional abort and may not satisfy the requirements for setting any of the EAX bits. This may result in an EAX value of 0.

#### RTM Memory Ordering

A successful RTM commit causes all memory operations in the RTM region to appear to execute atomically. A successfully committed RTM region consisting of an XBEGIN followed by an XEND, even with no memory operations in the RTM region, has the same ordering semantics as a LOCK prefixed instruction.

The XBEGIN instruction does not have fencing semantics. However, if an RTM execution aborts, then all memory updates from within the RTM region are discarded and are not made visible to any other logical processor.

#### RTM-Enabled Debugger Support

By default, any debug exception inside an RTM region will cause a transactional abort and will redirect control flow to the fallback instruction address with architectural state recovered and bit 4 in EAX set. However, to allow software debuggers to intercept execution on debug exceptions, the RTM architecture provides additional capability.

If bit 11 of DR7 and bit 15 of the IA32\_DEBUGCTL\_MSR are both 1, any RTM abort due to a debug exception (#DB) or

breakpoint exception (#BP) causes execution to roll back and restart from the XBEGIN instruction instead of the fallback address. In this scenario, the EAX register will also be restored back to the point of the XBEGIN instruction.

#### Programming Considerations

Typical programmer-identified regions are expected to transactionally execute and commit successfully. However, Intel® TSX does not provide any such guarantee. A transactional execution may abort for many reasons. To take full advantage of the transactional capabilities, programmers should follow certain guidelines to increase the probability of their transactional execution committing successfully.

This section discusses various events that may cause transactional aborts. The architecture ensures that updates performed within a transaction that subsequently aborts execution will never become visible. Only committed transactional executions initiate an update to the architectural state. Transactional aborts never cause functional failures and only affect performance.

#### Instruction Based Considerations

Programmers can use any instruction safely inside a transaction (HLE or RTM) and can use transactions at any privilege level. However, some instructions will always abort the transactional execution and cause execution to seamlessly and safely transition to a non-transactional path.

Intel® TSX allows for most common instructions to be used inside transactions without causing aborts. The following operations inside a transaction do not typically cause an abort:

- Operations on the instruction pointer register, general purpose registers (GPRs) and the status flags (CF, OF, SF, PF, AF, and ZF); and
- Operations on XMM and YMM registers and the MXCSR register.

However, programmers must be careful when intermixing SSE and AVX operations inside a transactional region. Intermixing SSE instructions accessing XMM registers and AVX instructions accessing YMM registers may cause transactions to abort. Programmers may use REP/REPNE prefixed string operations inside transactions. However, long strings may cause aborts. Further, the use of CLD and STD instructions may cause aborts if they change the value of the DF flag. However, if DF is 1, the STD instruction will not cause an abort. Similarly, if DF is 0, then the CLD instruction will not cause an abort.

Instructions not enumerated here as causing abort when used inside a transaction will typically not cause a transaction to abort (examples include but are not limited to MFENCE, LFENCE, SFENCE, RDTSC, RDTSCP, etc.).

The following instructions will abort transactional execution on any implementation:

XABORT  
CPUID  
PAUSE

In addition, in some implementations, the following instructions may always cause transactional aborts. These instructions are not expected to be commonly used inside typical transactional regions. However, programmers must not rely on these instructions to force a transactional abort, since whether they cause transactional aborts is implementation dependent.

Operations on X87 and MMX architecture state. This includes all MMX and X87 instructions, including the FXRSTOR and FXSAVE instructions.

Update to non-status portion of EFLAGS: CLI, STI, POPFD, POPFQ, CLTS.

Instructions that update segment registers, debug registers and/or control registers: MOV to DS/ES/FS/GS/SS, POP DS/ES/FS/GS/SS, LDS, LES, LFS, LGS, LSS, SWAPGS, WRFSBASE, WRGSBASE, LGDT, SGDT, LIDT, SIDT, LLDT, SLDT, LTR, STR, Far CALL, Far JMP, Far RET, IRET, MOV to DRx, MOV to CR0/CR2/CR3/CR4/CR8 and LMSW.

Ring transitions: SYSENTER, SYSCALL, SYSEXIT, and SYSRET.

TLB and Cacheability control: CLFLUSH, INVD, WBINVD, INVLPG, INVPCID, and memory instructions with a non-temporal hint (MOVNTDQA, MOVNTDQ, MOVNTI, MOVNTPD, MOVNTPS, and MOVNTQ).

Processor state save: XSAVE, XSAVEOPT, and XRSTOR. Interrupts: INTn, INTO.

IO: IN, INS, REP INS, OUT, OUTS, REP OUTS and their variants.

VMX: VMPTRLD, VMPTRST, VMCLEAR, VMREAD, VMWRITE, VMCALL, VMLAUNCH, VMRESUME, VMXOFF, VMXON, INVEPT, and INVVPID.

SMX: GETSEC.

UD2, RSM, RDMSR, WRMSR, HLT, MONITOR, MWAIT, XSETBV, VZEROUPPER, MASKMOVQ, and V/MASKMOVDQU.

#### Runtime Considerations

In addition to the instruction-based considerations, runtime events may cause transactional execution to abort. These may be due to data access patterns or micro-architectural implementation features. The following list is not a comprehensive discussion of all abort causes.

Any fault or trap in a transaction that must be exposed to software will be suppressed. Transactional execution will abort and execution will transition to a non-transactional execution, as if the fault or trap had never occurred. If an exception is not masked, then that un-masked exception will result in a transactional abort and the state will appear as if the exception had never occurred.

Synchronous exception events (#DE, #OF, #NP, #SS, #GP, #BR, #UD, #AC, #XF, #PF, #NM, #TS, #MF, #DB, #BP/INT3) that occur during transactional execution may cause an execution not to commit transactionally, and require a non-transactional execution. These events are suppressed as if they had never occurred. With HLE, since the non-transactional code path is identical to the transactional code path, these events will typically re-appear when the instruction that caused the exception is re-executed non-transactionally, causing the associated synchronous events to be delivered appropriately in the non-transactional execution. Asynchronous events (NMI, SMI, INTR, IPI, PMI, etc.) occurring during transactional execution may cause the transactional execution to abort and transition to a non-transactional execution. The asynchronous events will be pended and handled after the transactional abort is processed.

Transactions only support write-back cacheable memory type operations. A transaction may always abort if the transaction includes operations on any other memory type. This includes instruction fetches to UC memory type.

Memory accesses within a transactional region may require the processor to set the Accessed and Dirty flags of the referenced page table entry. The behavior of how the processor handles this is implementation specific. Some implementations may allow the updates to these flags to become externally visible even if the transactional region subsequently aborts. Some Intel® TSX implementations may choose to abort the transactional execution if these flags need to be updated. Further, a processor's page-table walk may generate

accesses to its own transactionally written but uncommitted state. Some Intel® TSX implementations may choose to abort the execution of a transactional region in such situations. Regardless, the architecture ensures that, if the transactional region aborts, then the transactionally written state will not be made architecturally visible through the behavior of structures such as TLBs.

Executing self-modifying code transactionally may also cause transactional aborts. Programmers must continue to follow the Intel recommended guidelines for writing self-modifying and cross-modifying code even when employing HLE and RTM. While an implementation of RTM and HLE will typically provide sufficient resources for executing common transactional regions, implementation constraints and excessive sizes for transactional regions may cause a transactional execution to abort and transition to a non-transactional execution. The architecture provides no guarantee of the amount of resources available to do transactional execution and does not guarantee that a transactional execution will ever succeed.

Conflicting requests to a cache line accessed within a transactional region may prevent the transaction from executing successfully. For example, if logical processor P0 reads line A in a transactional region and another logical processor P1 writes line A (either inside or outside a transactional region) then logical processor P0 may abort if logical processor P1's write interferes with processor P0's ability to execute transactionally.

Similarly, if P0 writes line A in a transactional region and P1 reads or writes line A (either inside or outside a transactional region), then P0 may abort if P1's access to line A interferes with P0's ability to execute transactionally. In addition, other coherence traffic may at times appear as conflicting requests and may cause aborts. While these false conflicts may happen, they are expected to be uncommon. The conflict resolution policy to determine whether P0 or P1 aborts in the above scenarios is implementation specific.

Generic Transaction Execution Embodiments:

According to "ARCHITECTURES FOR TRANSACTIONAL MEMORY", a dissertation submitted to the Department of Computer Science and the Committee on Graduate Studies of Stanford University in partial fulfillment of the requirements for the Degree of Doctor of Philosophy, by Austen McDonald, June 2009, incorporated by reference herein in its entirety, fundamentally, there are three mechanisms needed to implement an atomic and isolated transactional region: versioning, conflict detection, and contention management.

To make a transactional code region appear atomic, all the modifications performed by that transactional code region must be stored and kept isolated from other transactions until commit time. The system does this by implementing a versioning policy. Two versioning paradigms exist: eager and lazy. An eager versioning system stores newly generated transactional values in place and stores previous memory values on the side, in what is called an undo-log. A lazy versioning system stores new values temporarily in what is called a write buffer, copying them to memory only on commit. In either system, the cache is used to optimize storage of new versions.

To ensure serializability between transactions, conflicts must be detected and resolved. The two systems, i.e., the eager and lazy versioning systems, detect conflicts by implementing a conflict detection policy, either optimistic or pessimistic. An optimistic system executes transactions in parallel, checking for conflicts only when a transaction commits. A pessimistic systems check for conflicts at each load and store.

Similar to versioning, conflict detection also uses the cache, marking each line as either part of the read-set, part of the write-set, or both. The two systems resolve conflicts by implementing a contention management policy. Many contention management policies exist, some are more appropriate for optimistic conflict detection and some are more appropriate for pessimistic. Described below are some example policies.

Since each transactional memory (TM) system needs both versioning detection and conflict detection, these options give rise to four distinct TM designs: Eager-Pessimistic (EP), Eager-Optimistic (EO), Lazy-Pessimistic (LP), and Lazy-Optimistic (LO). Table 2 briefly describes all four distinct TM designs.

FIGS. 1 and 2 depict an example of a multicore TM environment. FIG. 1 shows many TM-enabled CPUs (CPU1 114a, CPU2 114b, etc.) on one die 100, connected with an interconnect 122 under management of an interconnect control 120a, 120b. Each CPU 114a, 114b (also known as a Processor) may have a split cache consisting of an Instruction Cache 116a, 116b for caching instructions from memory to be executed and a Data Cache 118a, 118b with TM support for caching data (operands) of memory locations to be operated on by the CPU 114a, 114b. In an implementation, caches of multiple dies 100 are interconnected to support cache coherency between the caches of the multiple dies 100. In an implementation, a single cache, rather than the split cache is employed holding both instructions and data. In implementations, the CPU caches are one level of caching in a hierarchical cache structure. For example each die 100 may employ a shared cache 124 to be shared amongst all the CPUs 114a, 114b on the die 100. In another implementation, each die 100 may have access to a shared cache 124, shared amongst all the processors of all the dies 100.

FIG. 2 shows the details of an example transactional CPU 114, including additions to support TM. The transactional CPU 114 (processor) may include hardware for supporting Register Checkpoints 126 and special TM Registers 128. The transactional CPU cache may have the MESI bits 130, Tags 140 and Data 142 of a conventional cache but also, for example, R bits 132 showing a line has been read by the CPU 114 while executing a transaction and W bits 138 showing a line has been written-to by the CPU 114 while executing a transaction.

A key detail for programmers in any TM system is how non-transactional accesses interact with transactions. By design, transactional accesses are screened from each other using the mechanisms above. However, the interaction between a regular, non-transactional load with a transaction containing a new value for that address must still be considered. In addition, the interaction between a non-transactional store with a transaction that has read that address must also be explored. These are issues of the database concept isolation.

A TM system is said to implement strong isolation, sometimes called strong atomicity, when every non-transactional load and store acts like an atomic transaction. Therefore, non-transactional loads cannot see uncommitted data and non-transactional stores cause atomicity violations in any transactions that have read that address. A system where this is not the case is said to implement weak isolation, sometimes called weak atomicity.

Strong isolation is often more desirable than weak isolation due to the relative ease of conceptualization and implementation of strong isolation. Additionally, if a programmer has forgotten to surround some shared memory references with transactions, causing bugs, then with strong isolation, the programmer will often detect that oversight using a simple

debug interface because the programmer will see a non-transactional region causing atomicity violations. Also, programs written in one model may work differently on another model.

Further, strong isolation is often easier to support in hardware TM than weak isolation. With strong isolation, since the coherence protocol already manages load and store communication between processors, transactions can detect non-transactional loads and stores and act appropriately. To implement strong isolation in software Transactional Memory (TM), non-transactional code must be modified to include read- and write-barriers; potentially crippling performance. Although great effort has been expended to remove many un-needed barriers, such techniques are often complex and performance is typically far lower than that of HTMs.

TABLE 2

Transactional Memory Design Space			
		VERSIONING	
		Lazy	Eager
CONFLICT DETECTION	Optimistic	Storing updates in a write buffer; detecting conflicts at commit time.	Not practical: waiting to update memory until commit time but detecting conflicts at access time guarantees wasted work and provides no advantage
	Pessimistic	Storing updates in a write buffer; detecting conflicts at access time.	Updating memory, keeping old values in undo log; detecting conflicts at access time.

Table 2 illustrates the fundamental design space of transactional memory (versioning and conflict detection). Eager-Pessimistic (EP)

This first TM design described below is known as Eager-Pessimistic. An EP system stores its write-set “in place” (hence the name “eager”) and, to support rollback, stores the old values of overwritten lines in an “undo log”. Processors use the W 138 and R 132 cache bits to track read and write-sets and detect conflicts when receiving snooped load requests. Perhaps the most notable examples of EP systems in known literature are LogTM and UTM.

Beginning a transaction in an EP system is much like beginning a transaction in other systems: `tm_begin()` takes a register checkpoint, and initializes any status registers. An EP system also requires initializing the undo log, the details of which are dependent on the log format, but often involve initializing a log base pointer to a region of pre-allocated, thread-private memory, and clearing a log bounds register.

Versioning: In EP, due to the way eager versioning is designed to function, the MESI 130 state transitions (cache line indicators corresponding to Modified, Exclusive, Shared, and Invalid code states) are left mostly unchanged. Outside of a transaction, the MESI 130 state transitions are left completely unchanged. When reading a line inside a transaction, the standard coherence transitions apply (S (Shared)→S, I (Invalid)→S, or I→E (Exclusive)), issuing a load miss as needed, but the R bit 132 is also set. Likewise, writing a line applies the standard transitions (S→M, E→I, I→M), issuing a miss as needed, but also sets the W (Written) bit 138. The first time a line is written, the old version of the entire line is loaded then written to the undo log to preserve it in case the current transaction aborts. The newly written data is then stored “in-place,” over the old data.

Conflict Detection: Pessimistic conflict detection uses coherence messages exchanged on misses, or upgrades, to look for conflicts between transactions. When a read miss occurs within a transaction, other processors receive a load request; but they ignore the request if they do not have the needed line. If the other processors have the needed line non-speculatively or have the line R 132 (Read), they downgrade that line to S, and in certain cases issue a cache-to-cache transfer if they have the line in MESI’s M or E state. However, if the cache has the line W 138, then a conflict is detected between the two transactions and additional action(s) must be taken.

Similarly, when a transaction seeks to upgrade a line from shared to modified (on a first write), the transaction issues an exclusive load request, which is also used to detect conflicts. If a receiving cache has the line non-speculatively, then the line is invalidated, and in certain cases a cache-to-cache transfer (M or E states) is issued. But, if the line is R 132 or W 138, a conflict is detected.

Validation: Because conflict detection is performed on every load, a transaction always has exclusive access to its own write-set. Therefore, validation does not require any additional work.

Commit: Since eager versioning stores the new version of data items in place, the commit process simply clears the W 138 and R 132 bits and discards the undo log.

Abort: When a transaction rolls back, the original version of each cache line in the undo log must be restored, a process called “unrolling” or “applying” the log. This is done during `tm_discard()` and must be atomic with regard to other transactions. Specifically, the write-set must still be used to detect conflicts: this transaction has the only correct version of lines in its undo log, and requesting transactions must wait for the correct version to be restored from that log. Such a log can be applied using a hardware state machine or software abort handler.

Eager-Pessimistic has the characteristics of: Commit is simple and since it is in-place, very fast. Similarly, validation is a no-op. Pessimistic conflict detection detects conflicts early, thereby reducing the number of “doomed” transactions. For example, if two transactions are involved in a Write-After-Read dependency, then that dependency is detected immediately in pessimistic conflict detection. However, in optimistic conflict detection such conflicts are not detected until the writer commits.

Eager-Pessimistic also has the characteristics of: As described above, the first time a cache line is written, the old value must be written to the log, incurring extra cache accesses. Aborts are expensive as they require undoing the log. For each cache line in the log, a load must be issued, perhaps going as far as main memory before continuing to the next line. Pessimistic conflict detection also prevents certain serializable schedules from existing.

Additionally, because conflicts are handled as they occur, there is a potential for livelock and careful contention management mechanisms must be employed to guarantee forward progress.

Lazy-Optimistic (LO)

Another popular TM design is Lazy-Optimistic (LO), which stores its write-set in a “write buffer” or “redo log” and detects conflicts at commit time (still using the R and W bits).

Versioning: Just as in the EP system, the MESI protocol of the LO design is enforced outside of the transactions. Once inside a transaction, reading a line incurs the standard MESI transitions but also sets the R bit 132. Likewise, writing a line sets the W bit 138 of the line, but handling the MESI transitions of the LO design is different from that of the EP design.

## 15

First, with lazy versioning, the new versions of written data are stored in the cache hierarchy until commit while other transactions have access to old versions available in memory or other caches. To make available the old versions, dirty lines (M lines) must be evicted when first written by a transaction. Second, no upgrade misses are needed because of the optimistic conflict detection feature: if a transaction has a line in the S state, it can simply write to it and upgrade that line to an M state without communicating the changes with other transactions because conflict detection is done at commit time.

Conflict Detection and Validation: To validate a transaction and detect conflicts, LO communicates the addresses of speculatively modified lines to other transactions only when it is preparing to commit. On validation, the processor sends one, potentially large, network packet containing all the addresses in the write-set. Data is not sent, but left in the cache of the committer and marked dirty (M). To build this packet without searching the cache for lines marked W, a simple bit vector is used, called a "store buffer," with one bit per cache line to track these speculatively modified lines. Other transactions use this address packet to detect conflicts: if an address is found in the cache and the R 132 and/or W 138 bits are set, then a conflict is initiated. If the line is found but neither R132 nor W 138 is set, then the line is simply invalidated, which is similar to processing an exclusive load.

To support transaction atomicity, these address packets must be handled atomically, i.e., no two address packets may exist at once with the same addresses. In an LO system, this can be achieved by simply acquiring a global commit token before sending the address packet. However, a two-phase commit scheme could be employed by first sending out the address packet, collecting responses, enforcing an ordering protocol (perhaps oldest transaction first), and committing once all responses are satisfactory.

Commit: Once validation has occurred, commit needs no special treatment: simply clear W 138 and R 132 bits and the store buffer. The transaction's writes are already marked dirty in the cache and other caches' copies of these lines have been invalidated via the address packet. Other processors can then access the committed data through the regular coherence protocol.

Abort: Rollback is equally easy: because the write-set is contained within the local caches, these lines can be invalidated, then clear W 138 and R 132 bits and the store buffer. The store buffer allows W lines to be found to invalidate without the need to search the cache.

Lazy-Optimistic has the characteristics of: Aborts are very fast, requiring no additional loads or stores and making only local changes. More serializable schedules can exist than found in EP, which allows an LO system to more aggressively speculate that transactions are independent, which can yield higher performance. Finally, the late detection of conflicts can increase the likelihood of forward progress.

Lazy-Optimistic also has the characteristics of: Validation takes global communication time proportional to size of write set. Doomed transactions can waste work since conflicts are detected only at commit time.

Lazy-Pessimistic (LP)

Lazy-Pessimistic (LP) represents a third TM design option, sitting somewhere between EP and LO: storing newly written lines in a write buffer but detecting conflicts on a per access basis.

Versioning: Versioning is similar but not identical to that of LO: reading a line sets its R bit 132, writing a line sets its W bit 138, and a store buffer is used to track W lines in the cache. Also, dirty (M) lines must be evicted when first written by a transaction, just as in LO. However, since conflict detection is

## 16

pessimistic, load exclusives must be performed when upgrading a transactional line from I, S→M, which is unlike LO.

Conflict Detection: LP's conflict detection operates the same as EP's: using coherence messages to look for conflicts between transactions.

Validation: Like in EP, pessimistic conflict detection ensures that at any point, a running transaction has no conflicts with any other running transaction, so validation is a no-op.

Commit: Commit needs no special treatment: simply clear W 138 and R 132 bits and the store buffer, like in LO.

Abort: Rollback is also like that of LO: simply invalidate the write-set using the store buffer and clear the W 138 and R 132 bits and the store buffer.

Eager-Optimistic (EO)

The LP has the characteristics of: Like LO, aborts are very fast. Like EP, the use of pessimistic conflict detection reduces the number of "doomed" transactions. Like EP, some serializable schedules are not allowed and conflict detection must be performed on each cache miss.

The final combination of versioning and conflict detection is Eager-Optimistic (EO). EO may be a less than optimal choice for HTM systems: since new transactional versions are written in-place, other transactions have no choice but to notice conflicts as they occur (i.e., as cache misses occur). But since EO waits until commit time to detect conflicts, those transactions become "zombies," continuing to execute, wasting resources, yet are "doomed" to abort.

EO has proven to be useful in STMs and is implemented by Bartok-STM and McRT. A lazy versioning STM needs to check its write buffer on each read to ensure that it is reading the most recent value. Since the write buffer is not a hardware structure, this is expensive, hence the preference for write-in-place eager versioning. Additionally, since checking for conflicts is also expensive in an STM, optimistic conflict detection offers the advantage of performing this operation in bulk. Contention Management

How a transaction rolls back once the system has decided to abort that transaction has been described above, but, since a conflict involves two transactions, the topics of which transaction should abort, how that abort should be initiated, and when should the aborted transaction be retried need to be explored. These are topics that are addressed by Contention Management (CM), a key component of transactional memory. Described below are policies regarding how the systems initiate aborts and the various established methods of managing which transactions should abort in a conflict.

Contention Management Policies

A Contention Management (CM) Policy is a mechanism that determines which transaction involved in a conflict should abort and when the aborted transaction should be retried. For example, it is often the case that retrying an aborted transaction immediately does not lead to the best performance. Conversely, employing a back-off mechanism, which delays the retrying of an aborted transaction, can yield better performance. STMs first grappled with finding the best contention management policies and many of the policies outlined below were originally developed for STMs.

CM Policies draw on a number of measures to make decisions, including ages of the transactions, size of read- and write-sets, the number of previous aborts, etc. The combinations of measures to make such decisions are endless, but certain combinations are described below, roughly in order of increasing complexity.

To establish some nomenclature, first note that in a conflict there are two sides: the attacker and the defender. The attacker is the transaction requesting access to a shared memory location. In pessimistic conflict detection, the attacker is the trans-



action issuing the load or load exclusive. In optimistic, the attacker is the transaction attempting to validate. The defender in both cases is the transaction receiving the attacker's request.

An Aggressive CM Policy immediately and always retries either the attacker or the defender. In LO, Aggressive means that the attacker always wins, and so Aggressive is sometimes called committer wins. Such a policy was used for the earliest LO systems. In the case of EP, Aggressive can be either defender wins or attacker wins.

Restarting a conflicting transaction that will immediately experience another conflict is bound to waste work—namely interconnect bandwidth refilling cache misses. A Polite CM Policy employs exponential backoff (but linear could also be used) before restarting conflicts. To prevent starvation, a situation where a process does not have resources allocated to it by the scheduler, the exponential backoff greatly increases the odds of transaction success after some  $n$  retries.

Another approach to conflict resolution is to randomly abort the attacker or defender (a policy called Randomized). Such a policy may be combined with a randomized backoff scheme to avoid unneeded contention.

However, making random choices, when selecting a transaction to abort, can result in aborting transactions that have completed “a lot of work”, which can waste resources. To avoid such waste, the amount of work completed on the transaction can be taken into account when determining which transaction to abort. One measure of work could be a transaction's age. Other methods include Oldest, Bulk TM, Size Matters, Karma, and Polka. Oldest is a simple timestamp method that aborts the younger transaction in a conflict. Bulk TM uses this scheme. Size Matters is like Oldest but instead of transaction age, the number of read/written words is used as the priority, reverting to Oldest after a fixed number of aborts. Karma is similar, using the size of the write-set as priority. Rollback then proceeds after backing off a fixed amount of time. Aborted transactions keep their priorities after being aborted (hence the name Karma). Polka works like Karma but instead of backing off a predefined amount of time, it backs off exponentially more each time.

Since aborting wastes work, it is logical to argue that stalling an attacker until the defender has finished their transaction would lead to better performance. Unfortunately, such a simple scheme easily leads to deadlock.

Deadlock avoidance techniques can be used to solve this problem. Greedy uses two rules to avoid deadlock. The first rule is, if a first transaction, T1, has lower priority than a second transaction, T0, or if T1 is waiting for another transaction, then T1 aborts when conflicting with T0. The second rule is, if T1 has higher priority than T0 and is not waiting, then T0 waits until T1 commits, aborts, or starts waiting (in which case the first rule is applied). Greedy provides some guarantees about time bounds for executing a set of transactions. One EP design (LogTM) uses a CM policy similar to Greedy to achieve stalling with conservative deadlock avoidance.

Example MESI coherency rules provide for four possible states in which a cache line of a multiprocessor cache system may reside, M, E, S, and I, defined as follows:

Modified (M): The cache line is present only in the current cache, and is dirty; it has been modified from the value in main memory. The cache is required to write the data back to main memory at some time in the future, before permitting any other read of the (no longer valid) main memory state. The write-back changes the line to the Exclusive state.

Exclusive (E): The cache line is present only in the current cache, but is clean; it matches main memory. It may be

changed to the Shared state at any time, in response to a read request. Alternatively, it may be changed to the Modified state when writing to it.

Shared (S): Indicates that this cache line may be stored in other caches of the machine and is “clean”; it matches the main memory. The line may be discarded (changed to the Invalid state) at any time.

Invalid (I): Indicates that this cache line is invalid (unused).

TM coherency status indicators (R 132, W 138) may be provided for each cache line, in addition to, or encoded in the MESI coherency bits. An R 132 indicator indicates the current transaction has read from the data of the cache line, and a W 138 indicator indicates the current transaction has written to the data of the cache line.

In another aspect of TM design, a system is designed using transactional store buffers. U.S. Pat. No. 6,349,361 titled “Methods and Apparatus for Reordering and Renaming Memory References in a Multiprocessor Computer System,” filed Mar. 31, 2000 and incorporated by reference herein in its entirety, teaches a method for reordering and renaming memory references in a multiprocessor computer system having at least a first and a second processor. The first processor has a first private cache and a first buffer, and the second processor has a second private cache and a second buffer. The method includes the steps of, for each of a plurality of gated store requests received by the first processor to store a datum, exclusively acquiring a cache line that contains the datum by the first private cache, and storing the datum in the first buffer. Upon the first buffer receiving a load request from the first processor to load a particular datum, the particular datum is provided to the first processor from among the data stored in the first buffer based on an in-order sequence of load and store operations. Upon the first cache receiving a load request from the second cache for a given datum, an error condition is indicated and a current state of at least one of the processors is reset to an earlier state when the load request from the given datum corresponds to the data stored in the first buffer.

The main implementation components of one such transactional memory facility are a transaction-backup register file for holding pre-transaction GR (general register) content, a cache directory to track the cache lines accessed during the transaction, a store cache to buffer stores until the transaction ends, and firmware routines to perform various complex functions. In this section a detailed implementation is described.

#### IBM zEnterprise EC12 Enterprise Server Embodiment

The IBM zEnterprise EC12 enterprise server introduces transactional execution (TX) in transactional memory, and is described in part in a paper, “Transactional Memory Architecture and Implementation for IBM System z” of Proceedings Pages 25-36 presented at MICRO-45, 1-5 Dec. 2012, Vancouver, British Columbia, Canada, available from IEEE Computer Society Conference Publishing Services (CPS), which is incorporated by reference herein in its entirety.

Table 3 shows an example transaction. Transactions started with TBEGIN are not assured to ever successfully complete with TEND, since they can experience an aborting condition at every attempted execution, e.g., due to repeating conflicts with other CPUs. This requires that the program supports a fallback path to perform the same operation non-transactionally, e.g., by using traditional locking schemes. This puts significant burden on the programming and software verification teams, especially where the fallback path is not automatically generated by a reliable compiler.

19

TABLE 3

Example Transaction Code			
loop	LHI	R0,0	*initialize retry count=0
	TBEGINC		*begin transaction
	JNZ	abort	*go to abort code if CC1=0
	LT	R1, lock	*load and test the fallback lock
	JNZ	lckbzy	*branch if lock busy
	... perform operation ...		
lckbzy	TEND		*end transaction
	... ..		
abort	TABORT		*abort if lock busy; this
			*resumes after TBEGINC
	JO	fallback	*no retry if CC=3
	AHI	R0, 1	*increment retry count
	CIJNL	R0,6,	*give up after 6 attempts
fallback		fallback	
	PPA	R0, TX	*random delay based on retry count
	... potentially wait for lock to become free ...		
	J	loop	*jump back to retry
	... ..		
OBTAIN	lock		*using Compare&Swap
	... perform operation ...		
	RELEASE	lock	
	... ..		

The requirement of providing a fallback path for aborted Transaction Execution (TX) transactions can be onerous. Many transactions operating on shared data structures are expected to be short, touch only a few distinct memory locations, and use simple instructions only. For those transactions, the IBM zEnterprise EC12 introduces the concept of constrained transactions; under normal conditions, the CPU 114 assures that constrained transactions eventually end successfully, albeit without giving a strict limit on the number of necessary retries. A constrained transaction starts with a TBEGINC instruction and ends with a regular TEND. Implementing a task as a constrained or non-constrained transaction typically results in very comparable performance, but constrained transactions simplify software development by removing the need for a fallback path. IBM's Transactional Execution architecture is further described in *z/Architecture, Principles of Operation*, Tenth Edition, SA22-7832-09 published September 2012 from IBM, incorporated by reference herein in its entirety.

A constrained transaction starts with the TBEGINC instruction. A transaction initiated with TBEGINC must follow a list of programming constraints; otherwise the program takes a non-filterable constraint-violation interruption. Exemplary constraints may include, but not be limited to: the transaction can execute a maximum of 32 instructions, all instruction text must be within 256 consecutive bytes of memory; the transaction contains only forward-pointing relative branches (i.e., no loops or subroutine calls); the transaction can access a maximum of 4 aligned octowords (an octoword is 32 bytes) of memory; and restriction of the instruction-set to exclude complex instructions like decimal or floating-point operations. The constraints are chosen such that many common operations like doubly linked list-insert/delete operations can be performed, including the very powerful concept of atomic compare-and-swap targeting up to 4 aligned octowords. At the same time, the constraints were chosen conservatively such that future CPU implementations can assure transaction success without needing to adjust the constraints, since that would otherwise lead to software incompatibility.

TBEGINC mostly behaves like XBEGIN in Intel® TSX or TBEGIN on IBM's zEC12 servers, except that the floating-point register (FPR) control and the program interruption

20

filtering fields do not exist and the controls are considered to be zero. On a transaction abort, the instruction address is set back directly to the TBEGINC instead of to the instruction after, reflecting the immediate retry and absence of an abort path for constrained transactions.

Nested transactions are not allowed within constrained transactions, but if a TBEGINC occurs within a non-constrained transaction it is treated as opening a new non-constrained nesting level just like TBEGIN would. This can occur, e.g., if a non-constrained transaction calls a subroutine that uses a constrained transaction internally.

Since interruption filtering is implicitly off, all exceptions during a constrained transaction lead to an interruption into the operating system (OS). Eventual successful finishing of the transaction relies on the capability of the OS to page-in the at most 4 pages touched by any constrained transaction. The OS must also ensure time-slices long enough to allow the transaction to complete.

TABLE 4

Transaction Code Example			
TBEGINC			*begin constrained transaction
	... perform operation ...		
TEND			*end transaction

Table 4 shows the constrained-transactional implementation of the code in Table 3, assuming that the constrained transactions do not interact with other locking-based code. No lock testing is shown therefore, but could be added if constrained transactions and lock-based code were mixed.

When failure occurs repeatedly, software emulation is performed using millicode as part of system firmware. Advantageously, constrained transactions have desirable properties because of the burden removed from programmers.

The IBM zEnterprise EC12 processor introduced the transactional execution facility. The processor can decode 3 instructions per clock cycle; simple instructions are dispatched as single micro-ops, and more complex instructions are cracked into multiple micro-ops 232b. The micro-ops (Uops 232b, shown in FIG. 3) are written into a unified issue queue 216, from where they can be issued out-of-order. Up to two fixed-point, one floating-point, two load/store, and two branch instructions can execute every cycle. A Global Completion Table (GCT) 232 holds every micro-op and a transaction nesting depth (TND) 232a. The GCT 232 is written in-order at decode time, tracks the execution status of each micro-op, and completes instructions when all micro-ops 232b of the oldest instruction group have successfully executed.

The level 1 (L1) data cache 240 (FIG. 3) is a 96 KB (kilo-byte) 6-way associative cache with 256 byte cache-lines and 4 cycle use latency, coupled to a private 1 MB (mega-byte) 8-way associative 2nd-level (L2) data cache 268 (FIG. 3) with 7 cycles use-latency penalty for L1 misses. L1 cache 240 (FIG. 3) is the cache closest to a processor and Ln cache is a cache at the nth level of caching. Both L1 240 (FIG. 3) and L2 268 (FIG. 3) caches are store-through. Six cores on each central processor (CP) chip share a 48 MB 3rd-level store-in cache, and six CP chips are connected to an off-chip 384 MB 4th-level cache, packaged together on a glass ceramic multi-chip module (MCM). Up to 4 multi-chip modules (MCMs) can be connected to a coherent symmetric multi-processor (SMP) system with up to 144 cores (not all cores are available to run customer workload).

## 21

Coherency is managed with a variant of the MESI protocol. Cache-lines can be owned read-only (shared) or exclusive; the L1 **240** (FIG. 3) and L2 **268** (FIG. 3) are store-through and thus do not contain dirty lines. The L3 and L4 caches are store-in and track dirty states. Each cache is inclusive of all its connected lower level caches.

Coherency requests are called “cross interrogates” (XI) and are sent hierarchically from higher level to lower-level caches, and between the L4s. When one core misses the L1 **240** (FIG. 3) and L2 **268** (FIG. 3) and requests the cache line from its local L3, the L3 checks whether it owns the line, and if necessary sends an XI to the currently owning L2 **268** (FIG. 3)/L1 **240** (FIG. 3) under that L3 to ensure coherency, before it returns the cache line to the requestor. If the request also misses the L3, the L3 sends a request to the L4 which enforces coherency by sending XIs to all necessary L3s under that L4, and to the neighboring L4s. Then the L4 responds to the requesting L3 which forwards the response to the L2 **268** (FIG. 3)/L1 **240** (FIG. 3).

Note that due to the inclusivity rule of the cache hierarchy, sometimes cache lines are XI’ed from lower-level caches due to evictions on higher-level caches caused by associativity overflows from requests to other cache lines. These XIs can be called “LRU XIs”, where LRU stands for least recently used.

Making reference to yet another type of XI requests, Demote-XIs transition cache-ownership from exclusive into read-only state, and Exclusive-XIs transition cache ownership from exclusive into invalid state. Demote-XIs and Exclusive-XIs need a response back to the XI sender. The target cache can “accept” the XI, or send a “reject” response if it first needs to evict dirty data before accepting the XI. The L1 **240** (FIG. 3)/L2 **268** (FIG. 3) caches are store through, but may reject demote-XIs and exclusive XIs if they have stores in their store queues that need to be sent to L3 before downgrading the exclusive state. A rejected XI will be repeated by the sender. Read-only-XIs are sent to caches that own the line read-only; no response is needed for such XIs since they cannot be rejected. The details of the SMP protocol are similar to those described for the IBM z10 by P. Mak, C. Walters, and G. Strait, in “IBM System z10 processor cache subsystem microarchitecture”, IBM Journal of Research and Development, Vol 53:1, 2009, which is incorporated by reference herein in its entirety.

#### Transactional Instruction Execution

FIG. 3 depicts example components of an example CPU. The instruction decode unit (IDU) **208** keeps track of the current transaction nesting depth (TND) **212**. When the IDU **208** receives a TBEGIN instruction, the nesting depth is incremented, and conversely decremented on TEND instructions. The nesting depth is written into the GCT **232** for every dispatched instruction. When a TBEGIN or TEND is decoded on a speculative path that later gets flushed, the IDU’s **208** nesting depth is refreshed from the youngest GCT **232** entry that is not flushed. The transactional state is also written into the issue queue **216** for consumption by the execution units, mostly by the Load/Store Unit (LSU) **280**. The TBEGIN instruction may specify a transaction diagnostic block (TDB) for recording status information, should the transaction abort before reaching a TEND instruction.

Similar to the nesting depth, the IDU **208**/GCT **232** collaboratively track the access register/floating-point register (AR/FPR) modification masks through the transaction nest; the IDU **208** can place an abort request into the GCT **232** when an AR/FPR-modifying instruction is decoded and the modification mask blocks that. When the instruction becomes next-to-complete, completion is blocked and the transaction aborts. Other restricted instructions are handled similarly,

## 22

including TBEGIN if decoded while in a constrained transaction, or exceeding the maximum nesting depth.

An outermost TBEGIN is cracked into multiple micro-ops depending on the GR-Save-Mask; each micro-op **232b** will be executed by one of the two fixed point units (FXUs) **220** to save a pair of GRs **228** into a special transaction-backup register file **224**, that is used to later restore the GR **228** content in case of a transaction abort. Also the TBEGIN spawns micro-ops **232b** to perform an accessibility test for the TDB if one is specified; the address is saved in a special purpose register for later usage in the abort case. At the decoding of an outermost TBEGIN, the instruction address and the instruction text of the TBEGIN are also saved in special purpose registers for a potential abort processing later on.

TEND and NTSTG are single micro-op **232b** instructions; NTSTG (non-transactional store) is handled like a normal store except that it is marked as non-transactional in the issue queue so that the LSU **280** can treat it appropriately. TEND is a no-op at execution time, the ending of the transaction is performed when TEND completes.

As mentioned, instructions that are within a transaction are marked as such in the issue queue **216**, but otherwise execute mostly unchanged; the LSU **280** performs isolation tracking as described in the next section.

Since decoding is in-order, and since the IDU **208** keeps track of the current transactional state and writes it into the issue queue **216** along with every instruction from the transaction, execution of TBEGIN, TEND, and instructions before, within, and after the transaction can be performed out-of order. An effective address calculator **236** is included in the LSU **280**. It is even possible (though unlikely) that TEND is executed first, then the entire transaction, and lastly the TBEGIN executes. Program order is restored through the GCT **232** at completion time. The length of transactions is not limited by the size of the GCT **232**, since general purpose registers (GRs) **228** can be restored from the backup register file **224**.

During execution, the program event recording (PER) events are filtered based on the Event Suppression Control, and a PER TEND event is detected if enabled. Similarly, while in transactional mode, a pseudo-random generator may be causing the random aborts as enabled by the Transaction Diagnostics Control.

#### Tracking for Transactional Isolation

The Load/Store Unit tracks cache lines that were accessed during transactional execution, and triggers an abort if an XI from another CPU (or an LRU-XI) conflicts with the footprint. If the conflicting XI is an exclusive or demote XI, the LSU rejects the XI back to the L3 in the hope of finishing the transaction before the L3 repeats the XI. This “stiff-arming” is very efficient in highly contended transactions. In order to prevent hangs when two CPUs stiff-arm each other, a XI-reject counter is implemented, which triggers a transaction abort when a threshold is met.

The L1 cache directory **240** is traditionally implemented with static random access memories (SRAMs). For the transactional memory implementation, the valid bits **244** (64 rows×6 ways) of the directory have been moved into normal logic latches, and are supplemented with two more bits per cache line: the TX-read **248** and TX-dirty **252** bits.

The TX-read **248** bits are reset when a new outermost TBEGIN is decoded (which is interlocked against a prior still pending transaction). The TX-read **248** bit is set at execution time by every load instruction that is marked “transactional” in the issue queue. Note that this can lead to over-marking if speculative loads are executed, for example on a mispredicted

## 23

branch path. The alternative of setting the TX-read bit at load completion time was too expensive for silicon area, since multiple loads can complete at the same time, requiring many read-ports on the load-queue.

Stores execute the same way as in non-transactional mode, but a transaction mark is placed in the store queue (STQ) 260 entry of the store instruction. At write-back time, when the data from the STQ 260 is written into the L1 240, the TX-dirty 252 bit in the L1-directory 256 is set for the written cache line. Store write-back into the L1 240 occurs only after the store instruction has completed, and at most one store is written back per cycle. Before completion and write-back, loads can access the data from the STQ 260 by means of store-forwarding; after write-back, the CPU 114 (FIG. 2) can access the speculatively updated data in the L1 240. If the transaction ends successfully, the TX-dirty bits 252 of all cache-lines are cleared, and also the TX-marks of not yet written stores are cleared in the STQ 260, effectively turning the pending stores into normal stores.

On a transaction abort, all pending transactional stores are invalidated from the STQ 260, even those already completed. All cache lines that were modified by the transaction in the L1 240, that is, have the TX-dirty bit 252 on, have their valid bits turned off, effectively removing them from the L1 240 cache instantaneously.

The architecture requires that before completing a new instruction, the isolation of the transaction read- and write-set is maintained. This isolation is ensured by stalling instruction completion at appropriate times when XIs are pending; speculative out-of order execution is allowed, optimistically assuming that the pending XIs are to different addresses and not actually cause a transaction conflict. This design fits very naturally with the XI-vs-completion interlocks that are implemented on prior systems to ensure the strong memory ordering that the architecture requires.

When the L1 240 receives an XI, L1 240 accesses the directory to check validity of the XI'd address in the L1 240, and if the TX-read bit 248 is active on the XI'd line and the XI is not rejected, the LSU 280 triggers an abort. When a cache line with active TX-read bit 248 is LRU'd from the L1 240, a special LRU-extension vector remembers for each of the 64 rows of the L1 240 that a TX-read line existed on that row. Since no precise address tracking exists for the LRU extensions, any non-rejected XI that hits a valid extension row the LSU 280 triggers an abort. Providing the LRU-extension effectively increases the read footprint capability from the L1-size to the L2-size and associativity, provided no conflicts with other CPUs 114 (FIG. 2) against the non-precise LRU-extension tracking causes aborts.

The store footprint is limited by the store cache size (the store cache is discussed in more detail below) and thus implicitly by the L2 size and associativity. No LRU-extension action needs to be performed when a TX-dirty cache line is LRU'd from the L1.

#### Store Cache

In prior systems, since the L1 240 and L2 268 are store-through caches, every store instruction causes an L3 store access; with now 6 cores per L3 and further improved performance of each core, the store rate for the L3 (and to a lesser extent for the L2) becomes problematic for certain workloads. In order to avoid store queuing delays a gathering store cache had to be added, that combines stores to neighboring addresses before sending them to the L3.

For transactional memory performance, it is acceptable to invalidate every TX-dirty cache line from the L1 240 on transaction aborts, because the L2 cache 268 is very close (7 cycles L1 miss penalty) to bring back the clean lines. How-

## 24

ever, it would be unacceptable for performance (and silicon area for tracking) to have transactional stores write the L2 268 before the transaction ends and then invalidate all dirty L2 cache lines on abort (or even worse on the shared L3).

The two problems of store bandwidth and transactional memory store handling can both be addressed with the gathering store cache. The cache 264 is a circular queue of 64 entries, each entry holding 128 bytes of data with byte-precise valid bits. In non-transactional operation, when a store is received from the LSU 280, the store cache 264 checks whether an entry exists for the same address, and if so gathers the new store into the existing entry. If no entry exists, a new entry is written into the queue, and if the number of free entries falls under a threshold, the oldest entries are written back to the L2 268 and L3 caches.

When a new outermost transaction begins, all existing entries in the store cache are marked closed so that no new stores can be gathered into them, and eviction of those entries to L2 268 and L3 is started. From that point on, the transactional stores coming out of the LSU 280 STQ 260 allocate new entries, or gather into existing transactional entries. The write-back of those stores into L2 268 and L3 is blocked, until the transaction ends successfully; at that point subsequent (post-transaction) stores can continue to gather into existing entries, until the next transaction closes those entries again.

The store cache 264 is queried on every exclusive or demote XI, and causes an XI reject if the XI compares to any active entry. If the core is not completing further instructions while continuously rejecting XIs, the transaction is aborted at a certain threshold to avoid hangs.

The LSU 280 requests a transaction abort when the store cache 264 overflows. The LSU 280 detects this condition when it tries to send a new store that cannot merge into an existing entry, and the entire store cache 264 is filled with stores from the current transaction. The store cache 264 is managed as a subset of the L2 268: while transactionally dirty lines can be evicted from the L1 240, they have to stay resident in the L2 268 throughout the transaction. The maximum store footprint is thus limited to the store cache size of 64×128 bytes, and it is also limited by the associativity of the L2 268. Since the L2 268 is 8-way associative and has 512 rows, it is typically large enough to not cause transaction aborts.

If a transaction aborts, the store cache 264 is notified and all entries holding transactional data are invalidated. The store cache 264 also has a mark per doubleword (8 bytes) whether the entry was written by a NTSTG instruction—those doublewords stay valid across transaction aborts.

#### Millicode-Implemented Functions

Traditionally, IBM mainframe server processors contain a layer of firmware called millicode which performs complex functions like certain CISC instruction executions, interruption handling, system synchronization, and RAS. Millicode includes machine dependent instructions as well as instructions of the instruction set architecture (ISA) that are fetched and executed from memory similarly to instructions of application programs and the operating system (OS). Firmware resides in a restricted area of main memory that customer programs cannot access. When hardware detects a situation that needs to invoke millicode, the instruction fetching unit 204 switches into "millicode mode" and starts fetching at the appropriate location in the millicode memory area. Millicode may be fetched and executed in the same way as instructions of the instruction set architecture (ISA), and may include ISA instructions.

For transactional memory, millicode is involved in various complex situations. Every transaction abort invokes a dedi-

25

cated millicode sub-routine to perform the necessary abort operations. The transaction-abort millicode starts by reading special-purpose registers (SPRs) holding the hardware internal abort reason, potential exception reasons, and the aborted instruction address, which millicode then uses to store a TDB if one is specified. The TBEGIN instruction text is loaded from an SPR to obtain the GR-save-mask, which is needed for millicode to know which GRs **228** to restore.

The CPU **114** (FIG. 2) supports a special millicode-only instruction to read out the backup-GRs and copy them into the main GRs. The TBEGIN instruction address is also loaded from an SPR to set the new instruction address in the Program Status Word (PSW) to continue execution after the TBEGIN once the millicode abort sub-routine finishes. That PSW may later be saved as program-old PSW in case the abort is caused by a non-filtered program interruption.

The TABORT instruction may be millicode implemented; when the IDU **208** decodes TABORT, it instructs the instruction fetch unit to branch into TABORT's millicode, from which millicode branches into the common abort sub-routine.

The Extract Transaction Nesting Depth (ETND) instruction may also be millicoded, since it is not performance critical; millicode loads the current nesting depth out of a special hardware register and places it into a GR **228**. The PPA instruction is millicoded; it performs the optimal delay based on the current abort count provided by software as an operand to PPA, and also based on other hardware internal state.

For constrained transactions, millicode may keep track of the number of aborts. The counter is reset to "0" on successful TEND completion, or if an interruption into the OS occurs (since it is not known if or when the OS will return to the program). Depending on the current abort count, millicode can invoke certain mechanisms to improve the chance of success for the subsequent transaction retry. The mechanisms involve, for example, successively increasing random delays between retries, and reducing the amount of speculative execution to avoid encountering aborts caused by speculative accesses to data that the transaction is not actually using. As a last resort, millicode can broadcast to other CPUs to stop all conflicting work, retry the local transaction, before releasing the other CPUs to continue normal processing. Multiple CPUs must be coordinated to not cause deadlocks, so some serialization between millicode instances on different CPUs is required.

Referring now to FIG. 4 an exemplary embodiment is illustrated for processing one or more instructions in a transaction based upon run-time conditions. In current technology, throughput gains may not scale as processors are added and the degree of parallelism increases. Throughput may slow to the point where contention for system serialization mechanisms, i.e., locks and semaphores, may even cause throughput to decline. Transactional memory may provide one approach to improving throughput. However, under some circumstances, such as where an accessed memory location (i.e., read usages) is modified by another thread (or processor core), a transaction can no longer maintain atomicity and will abort. Similarly, if any pending memory updates (i.e., write usages) are forced to be made visible to another thread (or processor core) or are backed-off due to cache subsystem conflicts before a transaction can be completed, the transaction will abort. It is therefore essential to track the memory accesses within a transaction that requires atomicity, in order to detect conflicts and potentially causes aborts. If the cache subsystem can no longer track a particular memory access for potential conflicts, e.g. when the transaction footprint of its memory accesses is larger than that supported by the hardware,

26

ware, the transaction will also abort. The transaction footprint tracked may include, for example, the memory addresses accessed for all instructions within the transaction that require atomicity. The success rate of a transaction's execution may be improved if the transaction's footprint may be determined with a greater degree of granularity. For example, a computing system that provides traditional transactional memory support may now conditionally suppress or enable tracking of memory accesses as part of the transaction footprint for a number of instructions while inside a transaction. The conditions, e.g. current condition code, architected states, or values at some memory location, can be checked at runtime of a transaction to override the default read and store footprint tracking policy, allowing an override mode where such memory usages may be suppressed from or enabled for transaction footprint tracking.

In an embodiment, a transaction override instruction is executed to cause the processor to be in transaction override mode (state). In the transaction override mode, memory accesses that would ordinarily be tracked for memory conflict in the transaction, may not be tracked, based on current processor state specified by the transaction override instruction. Thus memory accesses that would normally be managed as transaction atomic accesses may no longer be handled as requiring atomicity while in this override mode. In another embodiment, while in a transaction override mode, memory operands that ordinarily would not be tracked for memory conflicts in the transaction, may be tracked, based on the current processor state specified by the transaction override instruction. Thus memory accesses that would normally be assumed not requiring atomicity will now be treated as transaction atomic accesses while in this override mode.

A processor, such as CPU **114** (FIG. 2), may receive an instruction while executing within a transaction, **405**. If the instruction is received outside of a transaction, **405**, the CPU **114** (FIG. 2) will process the instruction in a traditional fashion; otherwise, the instruction will be processed in a transactional handling context. In general, a transaction may comprise memory accessing and non-memory accessing instructions. For example, although an arithmetic instruction such as MULTIPLY or ADD, may be within a transaction, these arithmetic instructions may not be memory-accessing (the operands may be in general registers rather than in memory), **410**. If the instruction received at **405** is accessing a cache line of a memory location, then the CPU **114** (FIG. 2) may determine whether to track the memory accessed by the instruction in the transaction footprint, based upon an architecturally defined tracking policy, **435**. For example, in one implementation transaction atomicity may be assumed across all memory accesses within a transaction and certain instructions or instruction prefixes may indicate the instruction operand is to be a non-transactional access, such as, for example, the NTSTG instruction mentioned earlier. In other architectural implementations, the default policy may be that transaction atomicity for an instruction within a transaction is not required unless explicitly specified. In both architectural implementations, the determination of non-default policy is statically determined by the usages of specific instructions or instruction prefixes in a program. When the architecturally defined default policy is not to track memory accesses in the transaction footprint, in order to include the memory access in the footprint the CPU **114** (FIG. 2) may override the default tracking policy with a the setting of an override mode. The override mode may be set by a special override instruction executed earlier, and may remain in effect for one or more instructions depending upon the specification, or until it is explicitly reset by another instruction. The special override

instruction may include further criteria for setting the override, such criteria may be based on current processor state conditions including, for example, current condition code values, values of a predefined set of bits in a Control Register (CR) or Special Purpose Register (SPR), or a dynamic value comparison to a set of bits in a specified general purpose, access, or floating point register. In one embodiment a special override instruction may specify a CR or SPR condition to be tested to override the default tracking policy. Thus, a bit-setting in a CR or SPR may enable non-default policy. The register's value can be loaded from memory at runtime with value previously updated by an operating system or another program working in conjunction with this transaction. After the override mode is set, a subsequent override reset instruction can be used to explicitly reset the override mode. In one embodiment the reset is conditional, for example, the override reset instruction may include further criteria for the reset based on current architected states, such as current condition code values, values of a predefined set of bits in a Control Register (CR) or Special Purpose Register (SPR), or a dynamic value comparison to a set of bits in a specified general purpose, access, or floating point register.

In the embodiment shown in Table 5, the override mode is specified by a prefix-type instruction, i.e. the mode is only set for the next sequential instruction (NSI) following the prefix-type instruction in program order, and then reset automatically

TABLE 5

Code Example Using Prefix Instruction	
CLIOVR R1,6,8	*compare immediate and override
ST R1, address	*influenced by default transactional handling override
.. perform operation ...	

In this example, using z-Architecture as reference, thus assuming an implementation where the default access handling is to add to the transaction footprint, an exemplary CLIOVR (Compare Immediate and Override) instruction performs both a comparison and override function. By comparing the value contained in the register R1 to "6", if they are equal, thus satisfying the mask value of 8, the override is enabled for the NSI. In this case, the override prevents the next instruction's storing access to be tracked in the transaction footprint. This prefix instruction allows a simple transactional store instruction to dynamically, based upon runtime conditions, (in this case whether the value of register R1 is "6") attain a similar conflict tracking footprint policy as the NTSTG (which was described earlier). If the compare would have not been satisfied, the default store conflict tracking is performed as part of the transaction footprint. The prefix-type instruction is illustrated here is a special case for understanding. An instruction prefix (or suffix) may be employed to provide the same function in an embodiment. Another variation of the prefix instruction can also be used whereby specifying the override mode will remain until a number of subsequent instructions are to be executed, or until reset by another instruction, or until the transaction ends for example.

If the tracking policy is overridden, **440**, then the memory access (usually tracked in the granularity of a cache line) may be added to the transaction footprint, **445**. However, it may be the case that the default tracking policy is not overridden, **440**, and as a result, the memory access is not added to the transaction footprint. Subsequently, the CPU **114** (FIG. 2) may reset from override mode back to the default tracking policy as specified by the special override instruction, **445**; such as

until next instruction is executed, until a number of instructions are executed, until reset by another instruction, or until the transaction ends.

If the architecturally defined default policy is to track memory accesses in the transaction footprint, **435**, then an override would have the effect of reversing the default tracking policy for the duration of the override mode, **450**. In that case, the memory access is not tracked in the transaction footprint. Alternatively, where the default account mode is unchanged, the memory access is added to the transaction footprint, **445**. Subsequently, the CPU **114** (FIG. 2) may reset from override mode back to the default tracking policy as specified by the special override instruction in various ways, for example, until next instruction is executed, until a number of instructions are executed, until reset by another instruction, or until the transaction ends.

Even though the instruction received at **405** is not accessing a memory, the CPU **114** (FIG. 2) may determine whether the detected instruction is one of the new or modified instructions that allow overriding the default tracking policy, **415**. For example, while a legacy LOAD instruction may not override the default tracking policy, an exemplary CLIOVR instruction as discussed earlier may indicate to the CPU **114** (FIG. 2) whether to set the override mode, based on state of the processor such as the value of a general register. If the CPU **114** (FIG. 2) determines that the override is to be done based on an inspection condition being satisfied, **425**, such as for example a Control Register (CR) setting, a condition code, or a value of a sets of bits in a GPR, then the override mode is set for the one or more instructions as specified in the special override instruction, **430**. However, if the inspection condition is not satisfied, no override will be done, **425**.

Referring now to FIG. 5, an exemplary embodiment is illustrated for processing an instruction in a transaction based upon run-time conditions. A processor, such as CPU **114** (FIG. 2), may receive an instruction while executing within a transaction, **505**. If the instruction is received outside of a transaction, the CPU **114** (FIG. 2) will process the instruction in a traditional fashion; otherwise, the instruction will be processed in a transactional handling context. In general, a transaction may comprise memory accessing and non-memory accessing instructions. In processing similar to that in **410** (FIG. 4), if the instruction received at **505** is accessing memory, **510**, then the CPU **114** (FIG. 2) may determine whether to track the memory accessed by the instruction in the transaction footprint, based upon the architecturally defined default policy, **535**. When the default is not to track the memory access in the transaction footprint, **515**, the CPU **114** (FIG. 2) may determine whether the detected instruction is one of the new or modified instructions that allow overriding the default tracking policy on its current memory access.

TABLE 6

Code Example Using Additional Mask	
CFI R1, 6	*compare immediate value and set condition code
NTOCSTG R5, 8, <address>	*conditionally enable non-default transactional handling
.. perform operation ...	

In the embodiment shown in Table 6, using the z-Architecture as the referenced architecture, a new instruction may influence its memory access as part of the transaction footprint tracking policy by specifying an additional 4-bits mask field. For example, in override mode, the exemplary NTOCSTG (Non Transactional On Condition Store) will suppress

29

the tracking in the transactional footprint if the current condition code is "0". In this code sequence, if the register R1 contains a value of "6", the prior Compare Immediate (CFI) instruction would set the condition code to "0", and thus satisfying the criteria set with the mask being a value of "8", dynamically enabling the NTOCSTG to perform like a NTSTG as described earlier. In this example, the runtime condition used for the override mode setting is the condition code value.

If the instruction can override default footprint tracking policy, 515, the override of operands of the instruction is further determined, 525 based on a runtime state condition, such as one or more bits in a CR, the condition code, or some other architected states. These conditions should be inspected to determine whether the override mode is to be performed for this instruction. Then upon successfully satisfying the inspection condition, 530, the memory access may be added to the transaction footprint, 565. In other embodiments, inspection may not be required to override the default tracking policy. In both embodiments, the architecturally defined default policy may be overridden only for the duration of the execution of this instruction, rather than for a plurality of instructions as in FIG. 4. Therefore, explicitly checking whether resetting back to the default account mode should be performed, as in 455 (FIG. 4), is not needed.

When the default tracking policy is to track the memory access in the transaction footprint, 535, then the CPU 114 (FIG. 2) may determine, as in 515, whether the detected instruction is one that allows overriding the default tracking policy (while accessing memory), 540. As in 525, if the CPU 114 (FIG. 2) determines that one or more architected states should be inspected prior to permitting the override mode for this instruction, 545, then upon successfully satisfying the inspection condition, the memory access may no longer be added to the transaction footprint, 550, bypassing 565. If, however, the detected instruction does not allow overriding the default tracking policy, 540, or if the detected instruction stated conditions are not satisfied, then the default account mode, i.e., to add the memory access to the transaction footprint, is performed, 565.

Referring now to FIG. 6, in an embodiment, dynamically adjusting a policy of whether a memory operand access of a memory accessing instruction is to be treated as a transaction atomic access in a transactional execution environment comprises: executing one or more processor-state-setting instructions, setting first state information in the processor, 610. The processor first state information is defined by an instruction set architecture (ISA) of the processor. After beginning execution of a transaction, processor second state information is set in the processor, 620. The processor second state information is defined by the ISA. Based on the processor second state information having a first value, 630, the processor enters transaction-footprint-override mode, 640. Based on the processor being in transaction-footprint-override mode, 650, memory conflicts are monitored according to the transaction-footprint-override mode state.

Referring now to FIG. 7, computing device 700 may include respective sets of internal components 800 and external components 900. Each of the sets of internal components 800 includes one or more processors 820; one or more computer-readable RAMs 822; one or more computer-readable ROMs 824 on one or more buses 826; one or more operating systems 828; one or more software applications executing the method of FIGS. 4-5; and one or more computer-readable tangible storage devices 830. The one or more operating systems are stored on one or more of the respective computer-readable tangible storage devices 830 for execution by one or

30

more of the respective processors 820 via one or more of the respective RAMs 822 (which typically include cache memory). In the embodiment illustrated in FIG. 7, each of the computer-readable tangible storage devices 830 is a magnetic disk storage device of an internal hard drive. Alternatively, each of the computer-readable tangible storage devices 830 is a semiconductor storage device such as ROM 824, EPROM, flash memory or any other computer-readable tangible storage device that can store a computer program and digital information.

Each set of internal components 800 also includes a R/W drive or interface 832 to read from and write to one or more computer-readable tangible storage devices 936 such as a thin provisioning storage device, CD-ROM, DVD, SSD, memory stick, magnetic tape, magnetic disk, optical disk or semiconductor storage device. The R/W drive or interface 832 may be used to load the device driver 840 firmware, software, or microcode to tangible storage device 936 to facilitate communication with components of computing device 700.

Each set of internal components 800 may also include network adapters (or switch port cards) or interfaces 836 such as a TCP/IP adapter cards, wireless WI-FI interface cards, or 3G or 4G wireless interface cards or other wired or wireless communication links. The operating system 828 that is associated with computing device 700, can be downloaded to computing device 700 from an external computer (e.g., server) via a network (for example, the Internet, a local area network or wide area network) and respective network adapters or interfaces 836. From the network adapters (or switch port adapters) or interfaces 836 and operating system 828 associated with computing device 700 are loaded into the respective hard drive 830 and network adapter 836. The network may comprise copper wires, optical fibers, wireless transmission, routers, firewalls, switches, gateway computers and/or edge servers.

Each of the sets of external components 900 can include a computer display monitor 920, a keyboard 930, and a computer mouse 934. External components 900 can also include touch screens, virtual keyboards, touch pads, pointing devices, and other human interface devices. Each of the sets of internal components 800 also includes device drivers 840 to interface to computer display monitor 920, keyboard 930 and computer mouse 934. The device drivers 840, R/W drive or interface 832 and network adapter or interface 836 comprise hardware and software (stored in storage device 830 and/or ROM 824).

Various embodiments of the disclosure may be implemented in a data processing system suitable for storing and/or executing program code that includes at least one processor coupled directly or indirectly to memory elements through a system bus. The memory elements include, for instance, local memory employed during actual execution of the program code, bulk storage, and cache memory which provide temporary storage of at least some program code in order to reduce the number of times code must be retrieved from bulk storage during execution. Input/Output or I/O devices (including, but not limited to, keyboards, displays, pointing devices, DASD, tape, CDs, DVDs, thumb drives and other memory media, etc.) can be coupled to the system either directly or through intervening I/O controllers. Network adapters may also be coupled to the system to enable the data processing system to become coupled to other data processing systems or remote printers or storage devices through intervening private or public networks. Modems, cable modems, and Ethernet cards are just a few of the available types of network adapters.

The present invention may be a system, a method, and/or a computer program product. The computer program product

31

may include a computer readable storage medium (or media) having computer readable program instructions thereon for causing a processor to carry out aspects of the present invention.

The computer readable storage medium can be a tangible device that can retain and store instructions for use by an instruction execution device. The computer readable storage medium may be, for example, but is not limited to, an electronic storage device, a magnetic storage device, an optical storage device, an electromagnetic storage device, a semiconductor storage device, or any suitable combination of the foregoing. A non-exhaustive list of more specific examples of the computer readable storage medium includes the following: a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), a static random access memory (SRAM), a portable compact disc read-only memory (CD-ROM), a digital versatile disk (DVD), a memory stick, a floppy disk, a mechanically encoded device such as punch-cards or raised structures in a groove having instructions recorded thereon, and any suitable combination of the foregoing. A computer readable storage medium, as used herein, is not to be construed as being transitory signals per se, such as radio waves or other freely propagating electromagnetic waves, electromagnetic waves propagating through a waveguide or other transmission media (e.g., light pulses passing through a fiber-optic cable), or electrical signals transmitted through a wire.

Computer readable program instructions described herein can be downloaded to respective computing/processing devices from a computer readable storage medium or to an external computer or external storage device via a network, for example, the Internet, a local area network, a wide area network and/or a wireless network. The network may comprise copper transmission cables, optical transmission fibers, wireless transmission, routers, firewalls, switches, gateway computers and/or edge servers. A network adapter card or network interface in each computing/processing device receives computer readable program instructions from the network and forwards the computer readable program instructions for storage in a computer readable storage medium within the respective computing/processing device.

Computer readable program instructions for carrying out operations of the present invention may be assembler instructions, instruction-set-architecture (ISA) instructions, machine instructions, machine dependent instructions, microcode, firmware instructions, state-setting data, or either source code or object code written in any combination of one or more programming languages, including an object oriented programming language, Smalltalk, C++ or the like, and conventional procedural programming languages, such as the "C" programming language or similar programming languages. The computer readable program instructions may execute entirely on the user's computer, partly on the user's computer, as a stand-alone software package, partly on the user's computer and partly on a remote computer or entirely on the remote computer or server. In the latter scenario, the remote computer may be connected to the user's computer through any type of network, including a local area network (LAN) or a wide area network (WAN), or the connection may be made to an external computer (for example, through the Internet using an Internet Service Provider). In some embodiments, electronic circuitry including, for example, programmable logic circuitry, field-programmable gate arrays (FPGA), or programmable logic arrays (PLA) may execute the computer readable program instructions by utilizing state information of the computer readable program instructions to

32

personalize the electronic circuitry, in order to perform aspects of the present invention.

Aspects of the present invention are described herein with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems), and computer program products according to embodiments of the invention. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer readable program instructions.

These computer readable program instructions may be provided to a processor of a general purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks. These computer readable program instructions may also be stored in a computer readable storage medium that can direct a computer, a programmable data processing apparatus, and/or other devices to function in a particular manner, such that the computer readable storage medium having instructions stored therein comprises an article of manufacture including instructions which implement aspects of the function/act specified in the flowchart and/or block diagram block or blocks.

The computer readable program instructions may also be loaded onto a computer, other programmable data processing apparatus, or other device to cause a series of operational steps to be performed on the computer, other programmable apparatus or other device to produce a computer implemented process, such that the instructions which execute on the computer, other programmable apparatus, or other device implement the functions/acts specified in the flowchart and/or block diagram block or blocks.

The flowchart and block diagrams in the Figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods, and computer program products according to various embodiments of the present invention. In this regard, each block in the flowchart or block diagrams may represent a module, segment, or portion of instructions, which comprises one or more executable instructions for implementing the specified logical function(s). In some alternative implementations, the functions noted in the block may occur out of the order noted in the figures. For example, two blocks shown in succession may, in fact, be executed substantially concurrently, or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved. It will also be noted that each block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, can be implemented by special purpose hardware-based systems that perform the specified functions or acts or carry out combinations of special purpose hardware and computer instructions.

Although preferred embodiments have been depicted and described in detail herein, it will be apparent to those skilled in the relevant art that various modifications, additions, substitutions and the like can be made without departing from the spirit of the disclosure, and these are, therefore, considered to be within the scope of the disclosure, as defined in the following claims.

What is claimed is:

1. A method for dynamically adjusting a policy of whether a memory operand access of a memory accessing instruction is to be treated as a transaction atomic access in a transactional execution environment, the method comprising:



33

executing, by a processor, one or more processor-state-setting instructions, the one or more processor-state-setting instructions causing processor first state information to be set in the processor, the processor first state information defined by an instruction set architecture (ISA) of the processor; and

beginning execution of a transaction, the execution comprising:

executing, by the processor, a transaction-policy-override instruction, the transaction-policy-override instruction causing processor second state information to be set in the processor, the processor second state information defined by the ISA of the processor;

based on the processor second state information having a first value, entering a transaction-footprint-override mode; and

based on the processor being in transaction-footprint-override mode, any one of ignoring memory conflicts that would otherwise be monitored and monitoring memory conflicts that would otherwise be ignored.

2. The method according to claim 1, wherein the transaction-policy-override instruction comprises any one of:

- an instruction preceding a memory accessing instruction;
- a prefix immediately preceding the memory accessing instruction; or
- a memory accessing instruction.

3. The method according to claim 1, wherein transaction memory accesses being treated as atomic comprise setting a read-set or a write-set indicator in a cache subsystem for indicating a corresponding cache line is to be monitored.

4. The method according to claim 1, further comprising:

determining whether a read access or a write access is to be treated as a transaction atomic access, based on an architecturally defined default policy, wherein an indicating program instruction overrides the architecturally defined default policy, based on the inspection of the one or more architected states being satisfied.

5. The method according to claim 1, wherein based on the processor second state information having a second value, exiting the transaction-footprint-override mode.

6. The method according to claim 4, wherein the overridden architecturally defined default policy is applied to a plurality of instructions based on at least one of:

- a completion of a specified number of instructions by the indicating program instruction;
- a resetting explicitly by another instruction; and
- an ending of the transaction.

7. A computer system for dynamically adjusting a policy of whether a memory operand access of a memory accessing instruction is to be treated as a transaction atomic access in a transactional execution environment, the computer system comprising:

- a memory; and
- a processor in communications with the memory, wherein the computer system is configured to perform a method, the method comprising:

executing, by the processor, one or more processor-state-setting instructions, the one or more processor-state-setting instructions causing processor first state information to be set in the processor, the processor first state information defined by an instruction set architecture (ISA) of the processor; and

beginning execution of a transaction, the execution comprising:

executing, by the processor, a transaction-policy-override instruction, the transaction-policy-override instruction causing processor second state information to be set in

34

the processor, the processor second state information defined by the ISA of the processor;

based on the processor second state information having a first value, entering a transaction-footprint-override mode; and

based on the processor being in transaction-footprint-override mode, any one of ignoring memory conflicts that would otherwise be monitored and monitoring memory conflicts that would otherwise be ignored.

8. The computer system according to claim 7, wherein the transaction-policy-override instruction comprises any one of:

- an instruction preceding a memory accessing instruction;
- a prefix immediately preceding the memory accessing instruction; or
- a memory accessing instruction.

9. The computer system according to claim 7, wherein transaction memory accesses being treated as atomic comprise setting a read-set or a write-set indicator in a cache subsystem for indicating a corresponding cache line is to be monitored.

10. The computer system according to claim 7, further comprising:

determining whether a read access or a write access is to be treated as a transaction atomic access, based on an architecturally defined default policy, wherein an indicating program instruction overrides the architecturally defined default policy, based on the inspection of the one or more architected states being satisfied.

11. The computer system according to claim 10, wherein based on the processor second state information having a second value, exiting the transaction-footprint-override mode.

12. The computer system according to claim 10, wherein the overridden architecturally defined default policy is applied to a plurality of instructions based on at least one of:

- a completion of a specified number of instructions by the indicating program instruction;
- a resetting explicitly by another instruction; and
- an ending of the transaction.

13. The computer system according to claim 12, wherein the overridden architecturally defined default policy is explicitly reset to the architecturally defined default policy by an instruction upon an inspection of one or more architected states.

14. A computer program product for dynamically adjusting a policy of whether a memory operand access of a memory accessing instruction is to be treated as a transaction atomic access in a transactional execution environment, comprising:

- a computer operating system embodied on a computer readable storage medium the computer operating system including program instructions executable by a processor, the program instructions comprising:

program instructions to execute, by the processor, one or more processor-state-setting instructions, the one or more processor-state-setting instructions causing processor first state information to be set in the processor, the processor first state information defined by an instruction set architecture (ISA) of the processor; and

program instruction to begin execution of a transaction, comprising:

program instructions to execute, by the processor, a transaction-policy-override instruction, the transaction-policy-override instruction causing processor second state information to be set in the processor, the processor second state information defined by the ISA of the processor;

35

program instructions to enter a transaction-footprint-over-  
ride mode based on the processor second state informa-  
tion having a first value; and

based on the processor being in transaction-footprint-over-  
ride mode, program instructions to perform any one of  
ignoring memory conflicts that would otherwise be  
monitored and monitoring memory conflicts that would  
otherwise be ignored.

**15.** The computer program product according to claim **14**,  
wherein the transaction-policy override instruction com-  
prises any one of:

an instruction preceding a memory accessing instruction;  
a prefix immediately preceding the memory accessing  
instruction; or  
a memory accessing instruction.

**16.** The computer program product according to claim **14**,  
wherein transaction memory accesses being treated as atomic  
accesses comprise program instructions to set a read-set or a  
write-set indicator in a cache subsystem for indicating a cor-  
responding cache line is to be monitored.

**17.** The computer program product according to claim **14**,  
further comprising:

program instructions to determine whether a read access or  
a write access is to be treated as a transaction atomic

36

access, based on an architecturally defined default  
policy, wherein an indicating program instruction over-  
rides the architecturally defined default policy, based on  
the inspection of the one or more architected states being  
satisfied.

**18.** The computer program product according to claim **17**,  
wherein based on the processor second state information  
having a second value, program instructions to exit the trans-  
action-footprint-override mode.

**19.** The computer program product according to claim **17**,  
wherein the overridden architecturally defined default policy  
is applied to a plurality of instructions based on at least one of:  
a completion of a specified number of instructions by the  
indicating program instruction;

a resetting explicitly by another instruction; and  
an ending of the transaction.

**20.** The computer program according to claim **19**, wherein  
the overridden architecturally defined default policy is explic-  
itly reset to the architecturally defined default policy by an  
instruction upon an inspection of one or more architected  
states.

\* \* \* \* \*